Modalities, Modes, and Models in Functional Neuroimaging

Karl J. Friston, *et al.*

Science **326**, 399 (2009);
DOI: 10.1126/science.1174521

22. L. Petreanu, T. Mao, S. M. Sternson, K. Svoboda, *Nature* **457**, 1142 (2009).
23. A. R. Adamantidis, F. Zhang, A. M. Aravanis, K. Deisseroth, L. de Lecea, *Nature* **450**, 420 (2007).
24. J. A. Cardin *et al.*, *Nature* **459**, 663 (2009).
25. V. S. Sohal, F. Zhang, O. Yizhar, K. Deisseroth, *Nature* **459**, 698 (2009).
26. P. Gorostiza, E. Y. Isacoff, *Science* **322**, 395 (2008).
27. A. Claridge-Chang *et al.*, *Cell*, 10.1016/j.cell.2009.08.034 (2009).
28. C. Wyart *et al.*, *Nature*, **461**, 407 (2009).
29. L. Luo, E. M. Callaway, K. Svoboda, *Neuron* **57**, 634 (2008).
30. A. Miyawaki, *Neuron* **48**, 189 (2005).
31. S. A. Hires, L. Tian, L. L. Looger, *Brain Cell Biol.* **36**, 69 (2008).
32. L. Sjulson, G. Miesenböck, *Physiology (Bethesda)* **22**, 47 (2007).
33. K. Ohki *et al.*, *Nature* **442**, 925 (2006).
34. M. Ng *et al.*, *Neuron* **36**, 463 (2002).
35. J. W. Wang, A. M. Wong, J. Flores, L. B. Vosshall, R. Axel, *Cell* **112**, 271 (2003).
36. Y. Shang, A. Claridge-Chang, L. Sjulson, M. Pypaert, G. Miesenböck, *Cell* **128**, 601 (2007).
37. L. Sjulson, G. Miesenböck, *Chem. Rev.* **108**, 1588 (2008).
38. K. Feldbauer *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 12317 (2009).
39. J. D. Clyne, G. Miesenböck, *Cell* **133**, 354 (2008).
40. R. A. Fisher, *J Min. Agric.* **33**, 503 (1926).
41. F. Rieke, D. Warland, R. de Ruyter van Steveninck, W. Bialek, *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, MA, 1997).
42. H. F. Judson, *The Eighth Day of Creation* (Simon and Schuster, New York, 1979).
43. I thank J. Flint, D. Kaetzel, P. Overton, E. Vrontou, and S. Waddell for comments. Work in my laboratory was supported by the Medical Research Council, the NIH, the Office of Naval Research, the Human Frontier Science Program, the Searle Scholars Program, and the Alfred P. Sloan, Beckman, Dana, Klingenstein, and McKnight foundations.

10.1126/science.1174520

REVIEW

Modalities, Modes, and Models in Functional Neuroimaging

Karl J. Friston

In this, the 21st century, human-brain mapping celebrates 21 years of cognitive activation studies. This review looks at imaging neuroscience and key ideas it has pursued; some ideas portend exciting developments, and others have failed gloriously. In terms of achievements, there is much to celebrate, in the sense that it is difficult to imagine modern neuroscience without brain imaging. I will look at recent advances from the perspectives of functional segregation and integration in the brain, paying special attention to approaches that deal with the distributed and integrated nature of neuronal processing and the questions they address.

Neuroimaging is now the predominant technique in behavioral and cognitive neuroscience. The volume of papers and number of fields it pervades are unrivaled (Fig. 1). Despite this, it is curiously difficult to summarize its achievements. The simplest summary falls back on the two guiding principles that shaped brain mapping at its inception: namely, functional segregation and integration. Neuroimaging has established functional segregation (the segregated or modular deployment of functional specialization within brain regions) as a fundament of brain organization. Furthermore, we can now characterize the integration of different brain areas in terms of functional and effective connectivity (Fig. 2). But beyond this, what have we learned? If you ask any imaging neuroscientist, they will recount exciting developments in their own field, ranging from the detailed functional architecture of retinotopically mapped visual cortex to the role of the ventral striatum in emotional learning. However, the question is more difficult to answer in terms of generic principles that underlie the brain's function and its relationship to anatomy. To see how people have tried to access these broader principles, I will look at recent trends in functional magnetic resonance imaging (fMRI), with a special focus on the questions that have been addressed. This focus is particularly important for functional neuroimaging, whose contributions will be measured by the depth of the questions asked, not the elegance of the method or, perhaps, the answers.

I first consider four themes that have caught people's imaginations recently and examine their underlying motivations, noting that there are many other exciting developments I could have addressed [such as genetics in neuroimaging, psychopharmacological fMRI, invasive and non-invasive electrophysiology, retinotopic mapping, computational anatomy, tractography with diffusion weighted imaging, lesion-deficit mapping, magnetic resonance spectroscopy, optical imaging, and technical advances such as polarization transfer (*1*)]. I then consider a couple of failures and conclude with a discussion of the implications for future directions; this discussion is illustrated with a few questions or model-led examples.

Multimodal Fusion

For years, we have heard about the promise of multimodal fusion, in which the spatial precision of fMRI will be complemented with the temporal precision of electroencephalography (EEG) to provide unprecedented spatiotemporal accuracy. However, this has not happened, despite the fact that we have the technology to acquire both modalities simultaneously (*2*) and have sophisticated biophysical models mapping from neuronal activity to both hemodynamic and electromagnetic measurements (*3*). So why is multimodal imaging not commonplace? Perhaps because there are many questions about functional anatomy that do not need bilateral spatial and temporal precision. Most questions about structure-function mapping and neuronal processing come in two flavors: where is it? or when is it? Functional magnetic resonance imaging is quite sufficient for questions of where and electromagnetic measurements [EEG and magnetoencephalography (MEG)] are the modalities of choice for questions of when; however, there are also questions about how imaging signals are generated that rest on fusion.

Multimodal fusion refers to the use of a common forward model of neuronal activity that explains different sorts of data. Several years ago, this was thought to be the best way to integrate fMRI and EEG because model parameters that are under-constrained by one modality might be informed by the other. In retrospect, this may have been a little misguided because the added value afforded by fusion requires unknown quantities generating data to express themselves in both modalities. Ironically, it may be that the complementary aspects of fMRI and EEG subvert the benefits of fusion. This may explain the success of simpler approaches to multimodal integration, in which the results from one modality constrain models of the other. These approaches use fMRI to provide precise spatial constraints (priors) on the source reconstruction of electromagnetic signals (*4*). Conversely, the temporal precision of EEG has been exploited in epilepsy research, in which explanatory variables based upon EEG features provide temporal constraints (in the form of explanatory variables or regressors) to model fMRI data (*5*).

So what questions call for fusion? A nice example is fusion of EEG and MEG data to measure evoked or induced responses, in which each modality alone is blind to certain sources (*6*). However, multimodal fusion really comes into its own when trying to understand the neurophysiology of brain-imaging signals and how they reflect underlying computations: Questions about func-

The Wellcome Trust Centre for Neuroimaging, University College London, Queen Square, London WC1N 3BG, UK. E-mail: k.friston@fil.ion.ucl.ac.uk

tional segregation are constrained by the resolution of fMRI. For example, a voxel (volume element of several mm^3) contains on average 5.5 million neurons, 10^{10} synapses, 22 km of dendrites, and 220 km of axons. Clearly, this precludes questions about functional specialization within the dendritic tree or indeed a cortical macrocolumn and calls for a broader notion of multimodal fusion, in which noninvasive imaging and invasive microscopic techniques are used to understand the principles of cortical computation that transcend spatial scales. I will return to this later but first consider techniques that try to find evidence for functional segregation at the voxel level.

Multivariate Pattern Classification, Decoding, and Mind-Reading

There has been immense interest recently in the use of multivariate pattern classification to infer the intentions or percepts of subjects by using fMRI measurements (mind-reading or decoding) (*7*–*9*). Conventional brain mapping tries to establish statistical dependencies between experimental manipulations and measured brain responses. This is related to a similar mapping in computational neuroscience, which refers to how neurons encode features in the outside world. The reverse mapping from measured physiological signals to the features encoded is called decoding. Decoding, reverse inference, or mind-reading uses multivariate analyses of fMRI data to classify the perceptual or cognitive state of a subject. Crucially, they harness patchy functional segregation (such as orientation selectivity in the visual cortex) at the voxel and subvoxel scale to search for patterns over voxels that best discriminate between experimental conditions.

Despite the allure of mind reading, one has to be wary of admiring the "emperor's new clothes." This is because multivariate pattern classification conflates multivariate with classification. Put briefly, their enhanced sensitivity and finessed characterizations of distributed responses rest on the use of multivariate models, not classification or reverse inference. Demonstrating a significant mapping between mental states and brain signals does not depend on the direction of the mapping (as with a significant correlation). In other words, showing that one can decode activity in the visual cortex to classify (above chance) a subject's percept is exactly the same as demonstrating significant visual cortex responses to perceptual changes. In this sense, all demonstrations of functionally specialized responses represent an implicit mind-reading. So what are the new questions behind decoding studies? If one looks closely, the questions are the same as in conventional encoding analyses: Basically, are there regionally specific correlates of some cognitive, perceptual, or sensorimotor state? However, when addressed with multivariate analyses the question pertains to distributed neuronal activity. This is because one is no longer testing for a dependency between an experimental variable and activity in one voxel but distributed responses over many voxels (Fig. 3). These distributed responses may have a fine-scale structure and be highly subject-specific, as opposed to the "blobs" identified in conventional mass-univariate analyses. This speaks to exciting developments, in which *t* tests used in conventional analyses (such as statistical parametric mapping) are replaced with statistics from multivariate models to provide maps of the mutual information between locally distributed cortical responses and experimental variables (*10*). Here, questions about functional segregation have not changed fundamentally but are framed in terms of distributed neural computations at the voxel or subvoxel scale. We will see below that multivariate models (such as eigenimage analysis and dynamic causal modeling) also play an important role in studying functional integration.

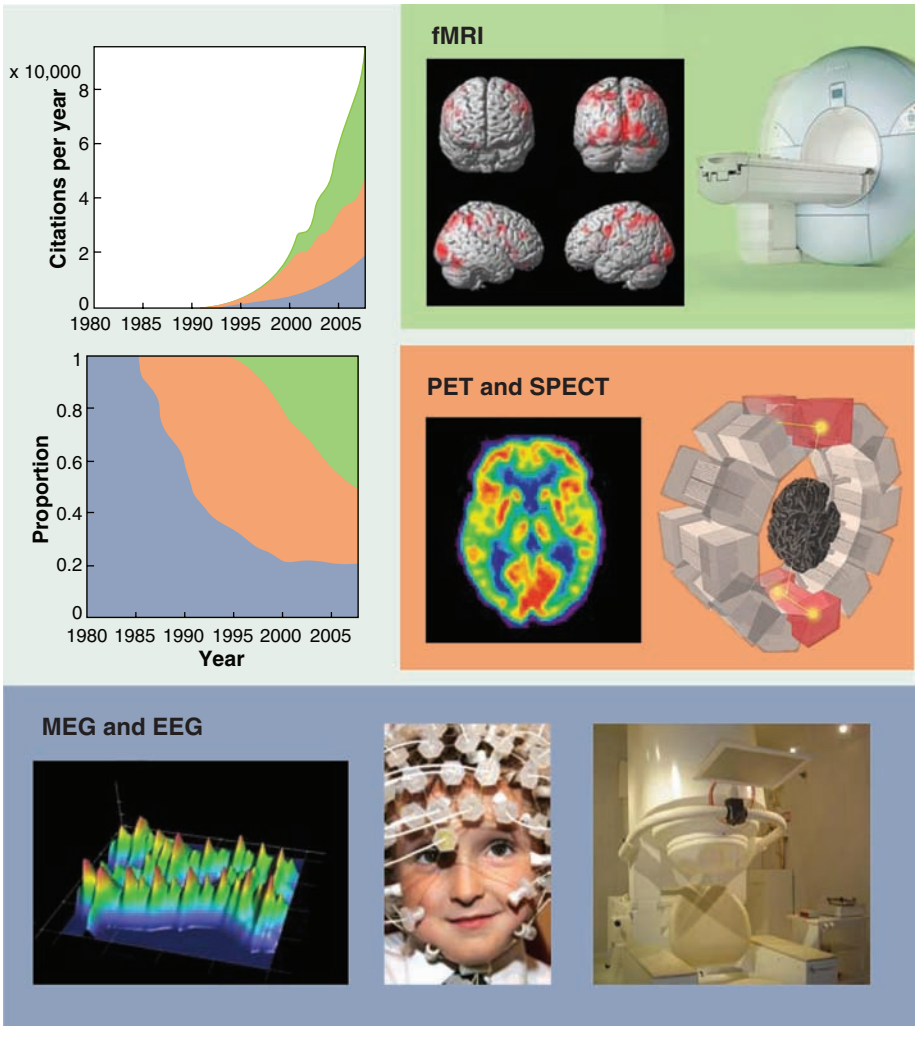

Fig. 1. Citation rates for the different modalities of functional neuroimaging. PET, positron emission tomography; SPECT, single-photon emission computed tomography. The citation rates [in 10,000 citations per year (**top left**) and proportion by modality (**middle left**)] are shown for the past 30 years. These data came from the ISI Web of Knowledge by searching for EEG OR MEG, PET OR SPECT, fMRI AND Brain with Topic = Neurosciences.

The Neurophysiological Basis of Imaging Signals

How many times have you read, "We know very little about the relationship between fMRI signals and their underlying neuronal causes"? In fact, decades of careful studies have clarified an enormous amount about the mapping between neuronal activity and hemodynamics (*11*–*14*). Furthermore, we know more than is sufficient to use fMRI for brain mapping. This is because the statistical models used to infer regionally specific responses make no assumptions about how neuronal responses are converted into measured signals (and that in particular do not assume this mapping is the same from voxel to voxel). In short, one does not need to know any neurophysiology to make precise inferences about where

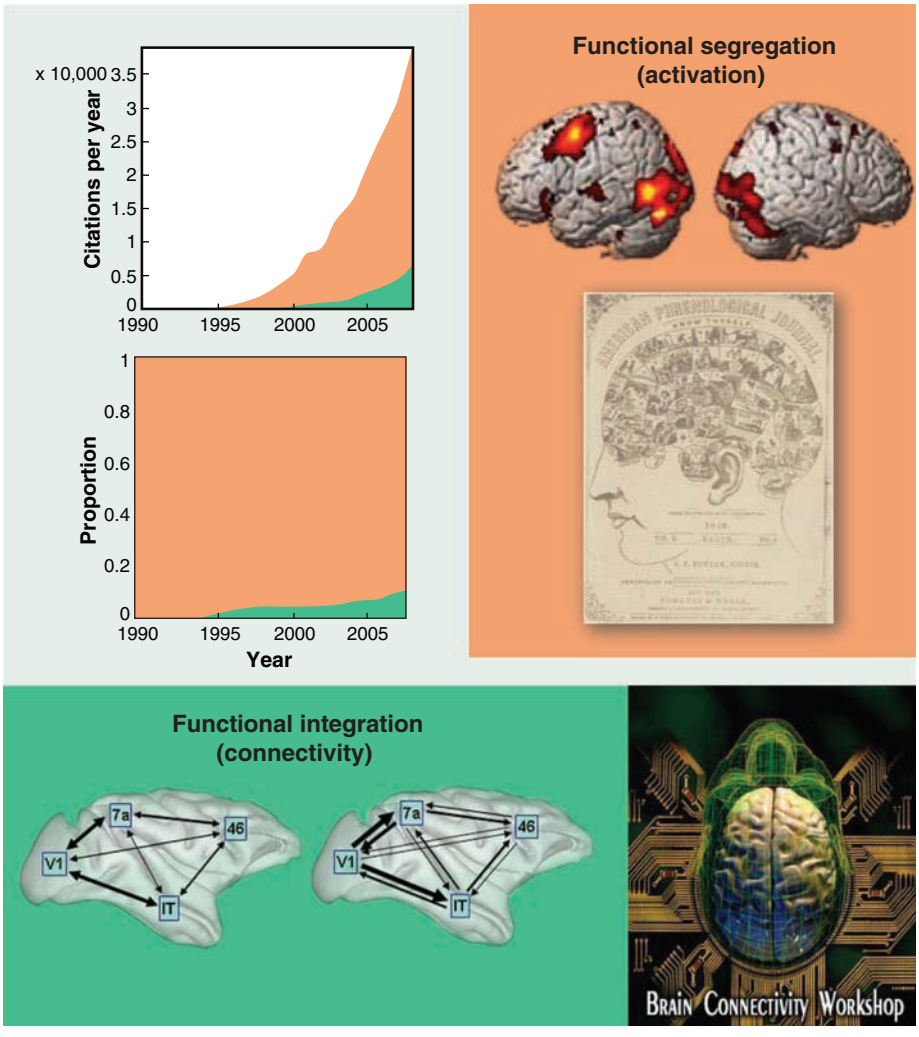

Fig. 2. Citation rates for the different applications of functional neuroimaging, derived in the same manner as for Fig. 1 but by searching for Activation AND functional imaging or Connectivity AND functional imaging with Topic = Neurosciences. This reflects the proportion of studies looking at functional segregation (activation) and those looking at integration (connectivity).

computations are taking place, provided there is some mapping between neuronal activity and hemodynamic responses (in a well-designed experiment, the only difference between conditions is neuronal activity, which is the only explanation for any hemodynamic difference).

So what are the questions asked when studying their relationship? These questions are about what can and cannot be inferred about neuronal activity from fMRI; for example, does fMRI reflect presynaptic inputs or postsynaptic firing (*14*); can it disambiguate between inhibitory and excitatory synaptic activity (*13*)? These questions address synaptic and microcircuit mechanisms in an attempt to disclose the relationship between spiking activity, local field potentials, and noninvasive imaging signals. Not only is this important for understanding what the brain is doing (in terms of local neuronal computations), it is critical for modeling distributed brain responses. This is because fMRI responses do not cause each other; they are caused by hidden neuronal states, and one needs to understand the mapping from neuronal states to measured signals to make sensible inferences about effective connectivity (the influence that one neuronal system has on another). For example, in fMRI, there is a distinction between models of effective connectivity used by dynamic causal modeling and econometric models (such as structural equation modeling and Granger causality). In dynamic causal modeling, neuronal states are modeled explicitly, allowing for regional variations in the hemodynamic response function. This variation violates the assumptions of econometric models. Recent multimodal studies in rat models of epilepsy have shown these variations can have a profound effect on inferences about effective connectivity (*15*). Another important example is the relationship between hemodynamics and the frequency of induced electrophysiological responses, in which higher frequencies (that may reflect neuromodulatory mechanisms) lead to increased fMRI signals (*16*).

Resting-State Correlations and Modes

There has been a recent upsurge in studies of fMRI signal correlations observed while the brain is at rest (*17*). These patterns seem to reflect anatomical connectivity (*18*, *19*) and can be characterized in terms of remarkably reproducible patterns of functional connectivity called modes (*20*). One of these modes recapitulates the pattern of deactivations observed in activation studies (the default mode) (*21*). These studies have been received with much excitement (see http://restingstate.stanford.edu/) and some ambivalence: On the one hand, they are very interesting. They suggest that, even at rest, endogenous activity in the brain is self-organizing and highly structured. On the other, their relationship to neuronal dynamics (*22*) and the questions they pose (*23*) are not always clear. One view of resting-state correlations is that they forego hypothesis testing because they preclude experimental manipulations. So why are they so interesting? Perhaps because they address functional integration and distributed processing and do so in the context of structure-function relationships. In this context, there are many mechanistic questions about the genesis of autonomous dynamics and the structures that support them. Some of the most interesting work in this field has come from computational anatomy and neuroscience. The emerging picture is that endogenous fluctuations are a consequence of dynamics on anatomical connectivity structures with particular scale-invariant and small-world characteristics (*24*, *25*). These are well-studied and universal characteristics of complex systems and suggest that we may be able to understand the brain in terms of universal phenomena. In short, endogenous fluctuations may be one way in which anatomy speaks to us through dynamics. Furthermore, they prompt important questions about how fluctuations shape evoked responses. In other words, are evoked brain responses and implicit computations affected by endogenous changes in its state (*26*)?

Some Interesting Failures

It is worthwhile to consider ideas that have not worked (and acknowledge those who tried to make them work). I will look at two examples, one from biophysics and one from neuroinformatics. A few years ago there was great excitement about the possibility of using fMRI to measure neuronal currents directly (*27*, *28*). Despite some impressive scientific prospecting, the hope that we will be able to measure neuronal activity directly on a millisecond and millimeter scale throughout the brain is now receding. Basically, the sorts of signals caused by neuronal activity currently appear to be too small to measure. However, it is comfort-

ing to reflect that we can already measure, with exquisite temporal precision, neuronal currents using EEG and MEG.

In 2000, there was a bold experiment to see if fMRI data sharing would provide added value for the imaging community (*29*) in terms of methodological cross-validation and the opportunity to reanalyze data. Despite laudable efforts, the experiment failed; why? Functional neuroimaging is a leader in providing informatics models for standardization (such as data formats and anatomical spaces) and software sharing (such as data analysis packages). However, unlike genomic or astrophysical data, fMRI data per se are very context sensitive. This means if data are optimal for your question, they are suboptimal for mine. This somewhat subverts the raison d'être for sharing. Furthermore, on a lighter note, a reviewer observed that "it's much too easy to collect … data (easier in fact than obtaining data from the data centre)."

Models and Questions

I have looked at recent developments in terms of the questions they address. The premise of this review is that the challenge for neuroimaging lies in specifying the questions or competing models (hypotheses) that it can explore. The implicit equivalence between questions and models may explain why many recent advances in characterizing brain-imaging data rest on model comparison. We have been comparing models with classical inference from the inception of brain mapping (comparing null and alternative models about regionally specific brain activations). However, the range and nature of models we could compare is growing rapidly. The conceptual challenge ahead may not lie in finessing the techniques at our disposal but informing the models used to explain data. I will consider four examples of modelled neuroimaging, all of which engage with fields beyond neuroimaging.

Computational Neuroscience

Perhaps the most obvious place to look for well-posed questions or models is theoretical neurobiology and computational neuroscience. There are many compelling examples that use fMRI to adjudicate among models of neuronal computations and their functional architectures. These studies rest on replacing traditional explanatory variables in statistical models of imaging data with quantities generated by computational models that are actually doing something. Nice examples here include the models of reinforcement and value learning (*30*–*33*). There is now an established tradition of taking a computational model that represents a hypothesis about how the brain evaluates contingencies, optimizing the model in relation to behavioral data, and using it to predict regionally specific fMRI responses (*33*). As a result of this approach, ventral striatal responses can now be treated as a proxy for unexpected rewards of the sort predicted by temporal difference models of learning. This is remarkable because a few years ago the only regionally specific correlates of reward came from invasive unit electrode recordings (*34*). Another nice example is in computational motor control, in which theories about optimal control can be evaluated against empirical fMRI data (*35*). The last example is the use of constructs from information theory to quantify novelty and surprise to see which parts of the brain encode causal regularities (or volatility) in our sensorium (*36*, *37*).

Neuroeconomics

A pleasing example of synergy between imaging and another field is the use of constructs from behavioral economics. Neuroimaging is now a primary modality for the study of neuroeconomics and allows researchers to establish the neuronal infrastructures that may be responsible for encoding and computing things like expected utility, discounting, and other variables that shape our choices in social or economic interactions (*38*, *39*). Another example, from the new field of social neuroscience, is the adoption of game theory to generate predictors of brain responses. In this context, model comparison allows one to lend a physiological validity to notions like guilt, regret, or altruism (*40*).

Biophysical Modeling

Model comparison also plays a key role in optimizing biophysical models. A pragmatic example here is the use of EEG data to adjudicate among forward models that map from sources in the brain to sensors. These models embody priors on the deployment of distributed neuronal sources and allow one to select the best prior assumptions. Put simply, one creates a number of models, each incorporating different but plausible assumptions (based on prior beliefs) about how data are generated. One then compares these models in terms of their evidence (the probability of obtaining data under that model). In some cases, this comparison is an implicit part of model optimization and can automatically switch off irrelevant neuronal sources (*41*). The same model comparison framework has proven extremely powerful in the comparison of dynamic causal models (*42*), in which hundreds of models (encoding different connectivity architectures) are

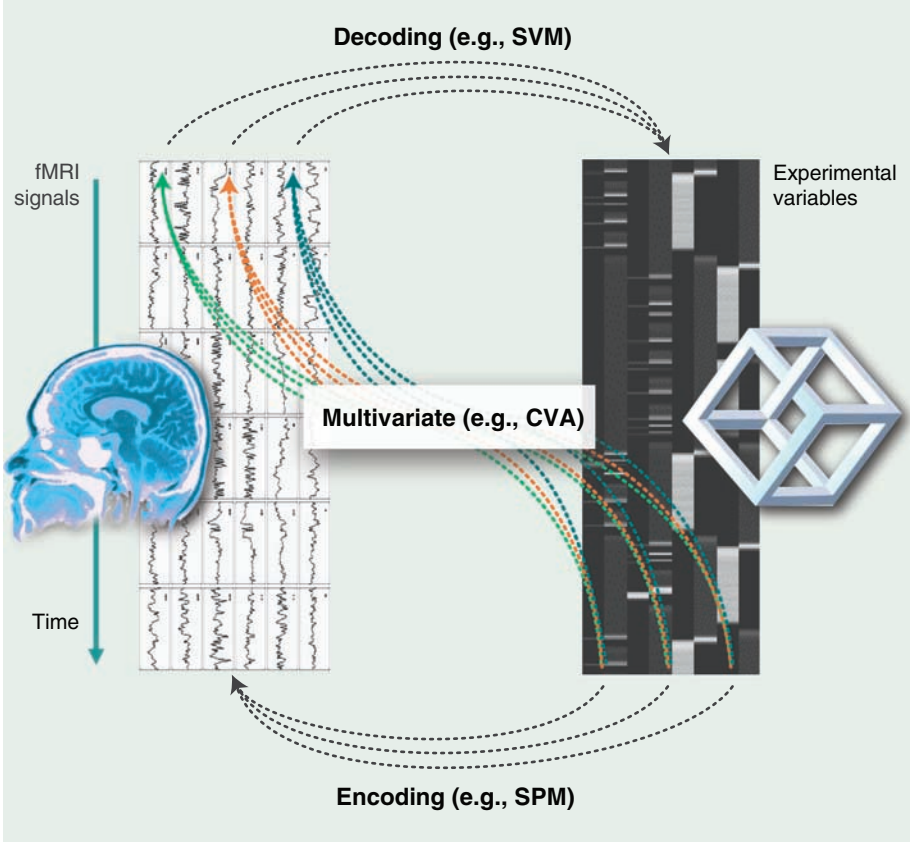

Fig. 3. Schematic showing the relationship between decoding (demonstrating a many-to-one mapping between activity in a local clique of voxels and an experimental variable), encoding (the conventional approach of showing a significant many-to-one mapping between experimental variables and activity in one voxel), and the more general multivariate many-to-many mapping between both sets of variables. SVM, support vector machine; SPM, statistical parametric mapping; CVA, canonical variates (or correlation) analysis.

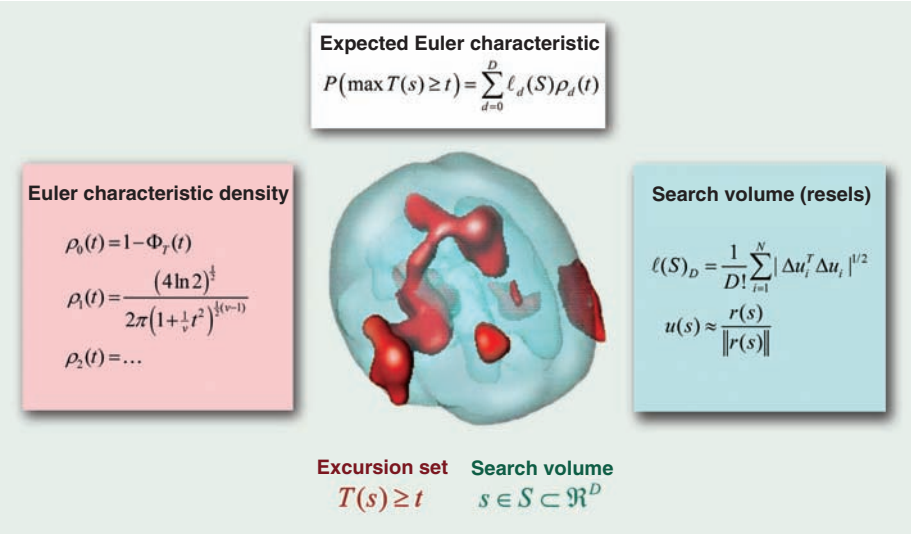

Fig. 4. Schematic showing some of the key equations behind topological inference. The basic idea is to split the problem of computing a *P* value (**top**) for a peak in a statistical image or map into two parts. The first part (**left**), the Euler characteristic density $\rho_D(t)$, depends only on the statistic and threshold t chosen for a D-dimensional search space; $s \in S$. This is the expected number of peaks per resolution element (resel). In this example, the equations pertain to the T statistic with a cumulative density Φ_T and ν degrees of freedom. The second part (**right**) is the search volume $\ell(S)_D$ measured in resels, which depends only on the shape and smoothness of the search space. This is defined in terms of the residuals $r(s)$ of the statistical test at each voxel; see (*44*) for details.

searched. These are an important class of models because many questions or hypotheses can be framed in terms of context-sensitive coupling between regions. Important questions here include the relative contribution of interhemispheric, bottom-up and top-down influences in cortical processing hierarchies and the implications for perception and action. These models have now been developed for fMRI, EEG, MEG, and local field potentials and are at the stage at which one can ask mechanistic questions about neuronal coupling by using, for example, animal models of Parkinson's disease (*43*).

Statistical Models

Imaging neuroscience has made some notable contributions to the physical sciences. In terms of statistical models for continuous data, neuroimaging has been at the forefront of developments, thanks to the contribution of people like Keith Worsley (who died prematurely a few months ago). Neuroimaging has essentially invented a new field of statistics, topological inference (based on random field theory), which underpins nearly all mainstream image analysis software (Fig. 4) (*44*). Similarly, some of the most advanced mathematical modeling of spatial data and their deformations has emerged in computational neuroanatomy (*45*). Finally, our systems models of connectivity and coupling are probably among the most developed in the biological sciences (*46*). These modeling contributions are not always heralded with the same applause as discovery science but underpin model comparison and the ability to ask questions of our data.

Conclusion

In summary, the most promising avenues for the future may rest on developing better models of our data that complement and exploit the richness of these data. These models may well already exist in other disciplines (such as machine learning, machine vision, computational neuroscience, and behavioral economics) and may enable the broader neurosciences to access neuroimaging so that key questions can be addressed in a theoretically grounded fashion. When I started writing this review, I was looking for headline themes and imminent breakthroughs. However, these may be less important than the vast number of incremental studies that pervade and enrich nearly every corner of neuroscience. Perhaps what we should be celebrating 21 years later is the fact that any bright young student or seasoned researcher can engage with an imaging unit and start to ask questions they are passionate about.

References and Notes

1. R. W. Adams *et al.*, *Science* **323**, 1708 (2009).
2. H. Laufs, J. Daunizeau, D. W. Carmichael, A. Kleinschmidt, *Neuroimage* **40**, 515 (2008).
3. P. A. Valdes-Sosa *et al.*, *Hum. Brain Mapp.* **30**, 2701 (2009)
4. A. M. Dale, E. Halgren, *Curr. Opin. Neurobiol.* **11**, 202 (2001).
5. L. Lemieux, *Neuroimaging Clin. N. Am.* **14**, 487 (2004).
6. S. Baillet, L. Garnero, G. Marin, J. P. Hugonin, *IEEE Trans. Biomed. Eng.* **46**, 522 (1999).
7. K. A. Norman, S. M. Polyn, G. J. Detre, J. V. Haxby, *Trends Cogn. Sci.* **10**, 424 (2006).
8. J. D. Haynes *et al.*, *Curr. Biol.* **17**, 323 (2007).
9. R. C. deCharms, *Nat. Rev. Neurosci.* **9**, 720 (2008).
10. R. R. Nandy, D. Cordes, *Magn. Reson. Med.* **50**, 354 (2003).
11. R. B. Buxton, E. C. Wong, L. R. Frank, *Magn. Reson. Med.* **39**, 855 (1998).
12. D. Attwell, C. Iadecola, *Trends Neurosci.* **25**, 621 (2002).
13. M. Lauritzen, L. Gold, *J. Neurosci.* **23**, 3972 (2003).
14. N. K. Logothetis, *Nature* **453**, 869 (2008).
15. O. David *et al.*, *PLoS Biol.* **6**, 2683 (2008).
16. J. B. Goense, N. K. Logothetis, *Curr. Biol.* **18**, 631 (2008).
17. B. Biswal, F. Z. Yetkin, V. M. Haughton, J. S. Hyde, *Magn. Reson. Med.* **34**, 537 (1995).
18. M. D. Greicius, K. Supekar, V. Menon, R. F. Dougherty, *Cereb. Cortex* **19**, 72 (2009).
19. A. K. Roy *et al.*, *Neuroimage* **45**, 614 (2009).
20. K. J. Friston, C. D. Frith, P. F. Liddle, R. S. Frackowiak, *J. Cereb. Blood Flow Metab.* **13**, 5 (1993).
21. M. E. Raichle *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 676 (2001).
22. R. M. Birn, K. Murphy, P. A. Bandettini, *Hum. Brain Mapp.* **29**, 740 (2008).
23. A. M. Morcom, P. C. Fletcher, *Neuroimage* **37**, 1073 (2007).
24. S. Achard, R. Salvador, B. Whitcher, J. Suckling, E. Bullmore, *J. Neurosci.* **26**, 63 (2006).
25. C. J. Honey, R. Kötter, M. Breakspear, O. Sporns, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 10240 (2007).
26. G. Hesselmann, C. A. Kell, A. Kleinschmidt, *J. Neurosci.* **28**, 14481 (2008).
27. J. Bodurka, P. A. Bandettini, *Magn. Reson. Med.* **47**, 1052 (2002).
28. L. Heller, B. E. Barrowes, J. S. George, *Hum. Brain Mapp.* **30**, 1 (2009).
29. J. D. Van Horn *et al.*, *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **356**, 1323 (2001).
30. J. P. O'Doherty, P. Dayan, K. Friston, H. Critchley, R. J. Dolan, *Neuron* **38**, 329 (2003).
31. P. F. Rodriguez, A. R. Aron, R. A. Poldrack, *Hum. Brain Mapp.* **27**, 306 (2006).
32. M. Haruno, M. Kawato, *Neural Netw.* **19**, 1242 (2006).
33. J. P. O'Doherty, A. Hampton, H. Kim, *Ann. N. Y. Acad. Sci.* **1104**, 35 (2007).
34. W. Schultz, *Neuron* **36**, 241 (2002).
35. S. T. Grafton, P. Schmitt, J. Van Horn, J. Diedrichsen, *Neuroimage* **39**, 1383 (2008).
36. T. E. Behrens, M. W. Woolrich, M. E. Walton, M. F. Rushworth, *Nat. Neurosci.* **10**, 1214 (2007).
37. B. A. Strange, A. Duggins, W. Penny, R. J. Dolan, K. J. Friston, *Neural Netw.* **18**, 225 (2005).
38. C. M. Kuhnen, B. Knutson, *Neuron* **47**, 763 (2005).
39. M. R. Delgado, A. Schotter, E. Y. Ozbay, E. A. Phelps, *Science* **321**, 1849 (2008).
40. I. Krajbich, R. Adolphs, D. Tranel, N. L. Denburg, C. F. Camerer, *J. Neurosci.* **29**, 2188 (2009).
41. D. Wipf, S. Nagarajan, *Neuroimage* **44**, 947 (2009).
42. W. D. Penny, K. E. Stephan, A. Mechelli, K. J. Friston, *Neuroimage* **22**, 1157 (2004).
43. R. J. Moran *et al.*, *Neuroimage* **42**, 272 (2008).
44. J. E. Taylor, K. J. Worsley, *J. Am. Stat. Assoc.* **102**, 913 (2007).
45. M. I. Miller, A. Qiu, *Neuroimage* **45** (suppl.), S16 (2009).
46. V. K. Jirsa, A. R. McIntosh, *Handbook of Brain Connectivity* (Springer, 2007). ISBN 3540714626, 9783540714620.

47. K.J.F. is funded by the Wellcome Trust. I would like to thank M. Bennett for preparing this manuscript and the reviewers for guidance in improving its presentation.

10.1126/science.1174521