

Chapter 2

Investigating Functional Cooperation in the Human Brain Using Simple Graph-Theoretic Methods

Michael L. Anderson, Joan Brumbaugh, and Aysu Şuben

Abstract This chapter introduces a very simple analytic method for mining large numbers of brain imaging experiments to discover functional cooperation between regions. We then report some preliminary results of its application, illustrate some of the many future projects in which we expect the technique will be of considerable use (including a way to relate fMRI to EEG), and describe a research resource for investigating functional cooperation in the cortex that will be made publicly available through the lab web site. One significant finding is that differences between cognitive domains appear to be attributable more to differences in patterns of cooperation between brain regions, rather than to differences in which brain regions are used in each domain. This is not a result that is predicted by prevailing localization-based and modular accounts of the organization of the cortex.

2.1 Introduction and Background

Hardly an issue of science or nature goes by without creating a stir over the discovery of “the” gene for some disease, trait, or predisposition, or “the” brain area responsible for some behavior or cognitive capacity. Of course, we know better; the isolable parts of complex systems like the brain or the human genome do what they do only in virtue of the cooperation of very many other parts, and often only by operating within and taking advantage of specific environmental and developmental

Michael L. Anderson

Department of Psychology, Franklin and Marshall College, Lancaster, PA 17604, USA; Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA, e-mail: michael . anderson@fandm . edu

Joan Brumbaugh

Department of Psychology, Franklin and Marshall College, Lancaster, PA 17604, USA

Aysu Şuben

Department of Psychology, Franklin and Marshall College, Lancaster, PA 17604, USA

contexts. But while it is true that we have gotten better about acknowledging the limitations of our instinctive reductionism – a bit of humility that the media would do well to absorb into its reporting – actual scientific practice has yet to be much affected by awareness of those limits. A recent case in point is John Anderson's project to map ACT-R components to brain regions [3]. The motivations for the project are of course entirely sound: if ACT-R is to be a realistic model of human cognition, then that model ought to have some significant, testable relationship to the neural bases of cognition. In this particular set of experiments, the authors identify eight ACT-R modules and match each one to a different region of interest. They then look for, and find, significant fit between the predictions for the BOLD signal in those regions, based on the activity of the ACT-R modules while solving a particular arithmetic task, and the measured BOLD signal in human participants performing the same task. On its face, this is an intriguing result and seems to offer compelling support for the ACT-R model. But the methodological assumption of the project – that there is a 1:1 mapping of ACT-R modules and brain areas – is highly suspect. Nor are the authors unaware of this difficulty, and in fact they specifically caution against making any inference from their approach to the functional organization of the brain:

Some qualifications need to be made to make it clear that we are not proposing a one-to-one mapping between these eight regions and the eight functions. First, other regions also serve these functions. Many areas are involved in vision and the fusiform gyrus has just proven to be the most useful to monitor. Similarly, many regions have been shown to be involved in retrieval, particularly the hippocampus. The prefrontal region is just the easiest to identify and seems to afford the best signal-to-noise ratio. Equally, we are not claiming these regions only serve one function. This paper has found some evidence for multiple functions. For instance, the motor regions are involved in rehearsal as well as external action (213–4).

Although we should appreciate the authors' candor here, the caveat seriously undermines the ability to interpret their results. If from the discovery that activity in an ACT-R module predicts the BOLD signal in specific brain region, we can neither infer that the region serves that specific function (because it is also activated in other tasks), nor that the function is served by that region (because other regions are activated by the target task), then we are not left with much. And yet despite the authors' awareness of these problems, they stick by the methodology that causes them.

Why might this be so? Naturally, all scientists are faced with the necessity of making simplifying abstractions to increase the tractability of their work; but as the authors found themselves, the assumption of a 1:1 mapping of modules to brain areas is not an approximation to reality, but appears to be fundamentally misleading. So what would account for the fact that they persist in applying methodological assumptions that they know to be inadequate? Given the scientific stature of the

authors, the question prompts reflection on the range and adequacy of the methodological tools actually available for work in this area. One sticks with improper tools only when the other options appear even worse. And while there are indeed more sophisticated tools for cooperation-sensitive investigations of neuroscientific data, those techniques are typically highly complex, hard to master, and – most importantly – produce results that can be difficult to interpret.

To help address these related problems, this chapter will describe a very simple analytical technique that we have been using in our lab to make cooperation-sensitive investigations tractable. In this chapter, we will outline that method, report some preliminary results of its application, and illustrate some of the many future projects in which we expect this technique (and the underlying database of brain imaging studies) will be of considerable use.

2.2 Graph Theory and Neuroscience

A graph is a set of objects called points, vertices, or nodes connected by links called lines or edges. Graphs have proven to be a convenient format to represent relationships in very many different areas, including computer networks, telephone calls, airline route maps, and social interactions [18, 19]. In neuroscience, graphs have been used for such purposes as investigating neural connectivity patterns [27], correcting brain images [17], and analyzing the patterns of neural activations in epilepsy [32]. Nevertheless graphs and graph theory – the branch of mathematics concerned with exploring the topological properties of graphs [15] – remain at this time underutilized tools with enormous potential to advance our understanding of the operations of the brain.

Our approach to investigating functional cooperation in the cortex involves building co-activation graphs, based on applying some simple data analysis techniques to large numbers of brain imaging studies. The method consists of two steps: first, choosing a spatial segmentation of the cortex to represent as nodes (current work uses Brodmann areas, but alternate segmentation schemes could easily be used; see below); and second, performing some simple analyses to discover which regions – which nodes – are statistically likely to be co-active. These relationships are represented as edges in our graphs.

For this second step we proceed in the following way. Given a database of brain imaging studies containing information about brain activations in various contexts (we describe the particular database we have been using in the next section), we first determine the chance likelihood of activation for each region by dividing the number of experiments in which it is reported to be active by the total number of experiments in the database. Then, for each pair of regions, we use a χ^2 measure to determine if the regions are more (or less) likely to be co-active than would be predicted by chance. We also perform a binomial analysis, since a binomial measure can provide directional information. (It is sometimes the case that, while area A and area B are co-active more (or less) often than would be predicted by chance, the

effect is asymmetric, such that area B is more active when area A is active, but not the reverse.)

Figure 2.1 shows the results of one such analysis, for a set of action and attention tasks. The graphs represent Brodmann areas that are significantly more likely than chance to be co-active ($\chi^2 > 3.84$); it is hypothesized that the network of co-activated areas revealed by such analysis represents those areas of the cortex that cooperate to perform the cognitive tasks in the given domain. The co-activation graphs are superimposed on an adjacency graph (where edges indicate that the Brodmann areas share a physical border in the brain) for ease of visual comparison.




Fig. 2.1: Cortex represented as adjacency + co-activation graphs. Here the Brodmann areas are nodes, with *black lines* between adjacent areas and *orange lines* between areas showing significant co-activation. The graph on the *left* shows co-activations from 56 action tasks, and the graph on the *right* shows co-activations from 77 attention tasks. Edges determined using the threshold $\chi^2 > 3.84$. Graphs rendered with aiSee v. 2.2.

Note that co-activation analysis is similar to, but distinct from, the approach adopted by [31] in discovering “functional connectivity.” The main difference is that edges in functional connectivity graphs indicate temporal co-variation between brain regions. Moreover, the results they report generally represent the dynamics of simulated neural networks (based on the structure of biological brain networks), rather than the static analysis of data-mining imaging experiments. Hence we adopt the term “functional cooperation” to distinguish our results from theirs. Nevertheless, there is presumably much to be gained by leveraging both sorts of analysis; in a later section we describe one such future project for bringing co-activation and co-variation graphs together.

The results of such analysis are not just visually striking, but afford the application of some well-understood mathematical techniques to better understand features of brain organization and functional cooperation. Of course, exactly what sorts of

techniques are appropriate, and how the end results should be interpreted, depend a great deal on the nature of the underlying data. Thus, in the next section we describe the database that we have been working with and how other researchers can get access to it for their own use. Then, in the final section, we will describe some of the projects to which we have applied this resource and some of the future possibilities.

2.3 A Database of Imaging Experiments

Over the last year or so we have compiled a database containing 665 experiments in 18 cognitive domains. The database currently consists of every qualifying imaging study in the *Journal of Cognitive Neuroscience* from 1996 to 2006, as well as the 135 experiments from [11] that were used in previous studies [4, 6]. To qualify for inclusion in the database, the study had to be conducted on healthy adults and to use a subtraction-based methodology for analyzing results. The database contains only post-subtraction activations. The data recorded for each experiment include the publication citation, the domain and sub-domain, the imaging method, the Talairach coordinates of each reported activation, the Brodmann area of each reported activation, the relative placement of the activation in the Brodmann area (e.g., frontal, caudal, ventral, dorsal), and the comparison used to generate the results. The domain labels are consistent with those used by the BrainMap database [22]. For experiments where coordinates were reported in MNI coordinates, a software package called GingerALE was used to translate these into Talairach coordinates [21]. When the authors of the study reported the Brodmann areas of their activations, these were recorded as reported. Where the authors did not report Brodmann areas, a software package called the Talairach demon [24] was used to provide Brodmann area labels for the coordinates. This program reports a distance in millimeters from the coordinate to the reported Brodmann area; this is the range, and it is recorded in cases where the BA label was generated using the software. The range is useful for excluding from analysis Brodmann area labels for coordinates that are further than desired from the reported area. Our plans are to continue to add to the database and analysis, and to publish versions at 1 year intervals beginning in the fall of 2008. The published versions of the database will contain the base data detailed above, as well as co-activation graphs, and will be prepared according to the following procedure: first, we will only include in the co-activation analysis sample domains containing some minimum number of experiments (e.g., 50 or 100, to be determined by what is feasible given the state of the database at that time). Having identified these domains, we will generate a concordance of authors to be sure that no individual labs are overrepresented in any given domain. The samples will be balanced by lab by randomly excluding experiments from overrepresented authors. At this point we will choose a target n based on the number of experiments in the domain containing the fewest number of experiments. An equal number of experiments will be

randomly selected from the remaining domains. This set of experiments, equally balanced between the domains, will be the sample for that year's co-activation analysis.

On this balanced sample we will run at least the following kinds of analysis. (1) For each domain, and for the entire set, we will generate a co-activation graph, constructed using the method outlined above, using Brodmann areas as nodes, and including only activations with a range (see above) of less than 5 mm. The calculated chance of activation and co-activation, as well as the binomial probability and χ^2 value will be reported for each pair of Brodmann areas, allowing researchers to set their own probability thresholds. (2) For each of the co-activation graphs, we will do a clique analysis (see below). Lancaster et al. [23] review some methods for generating cliques from brain activation data, and there are many other well-established methods for extracting cliques of various descriptions from graphs [1, 8, 9, 16]. Finally, (3) for all of the co-activation graphs and cliques, we will project them onto the adjacency graph (shown above) and calculate the average minimum graph distance (the "scatter" in the cortex) of the included nodes. All of this data will be made available for download from the lab web site, at

http://www.agcognition.org/brain_network

Before moving on to the next section, where we describe some of the uses to which these data have been put, and how it can be applied in the future, it is worth saying a word about our reliance on Brodmann areas as the basis for the analyses. It is of course legitimate to wonder whether the sub-division of the cortex into Brodmann areas will be a feature of our final functional map of the human brain; one rather suspects it will be fully superseded by some yet-to-be developed topographical scheme. Yet Brodmann areas remain the lingua franca in Cognitive Neuroscience for reporting findings, and sticking to this tradition will make results using these analyses easier to relate to past findings. Moreover, for the purposes we have described here – investigating the functional cooperation between brain areas involved in supporting different functions – virtually any consistent spatial division of the brain will do, and regions the size of Brodmann areas offer adequate spatial resolution for the required analysis. For, while the spatial resolution of a single fMRI image is on the order of 3 mm or better, there are questions both about the accuracy and precision of repeated fMRI, both within and between participants, effectively reducing its functional resolution [28]. It is arguable, then, that the use of Brodmann-sized regions of the cortex for representing the contribution of individual brain areas to cognitive tasks is consistent with the realistic (conservatively estimated) spatial resolution of current imaging technologies [10, 34]. In any case, it should be noted that the coordinates of each activation are also recorded in the database; if a Brodmann-based spatial scheme does not appear to produce useful or legitimate results, other spatial divisions of the cortex can certainly be substituted, and the very same sort of analysis performed. For instance, one can use the ALE (activation likelihood estimates) paradigm [33] to extract probable activations for arbitrarily defined neural volumes and build graphs from these data [23].

2.4 The Usefulness of Co-activation Graphs

With brain imaging data in this format, it becomes possible to formulate some very simple questions and use some well-understood methods to answer them. For instance, a long-standing project in our lab has been adjudicating between functional topographies of the brain based on the principle of localization and those based on the principle of redeployment. Localization-based approaches to functional topography, insofar as they typically expect brain regions to be dedicated to a small and domain-restricted set of cognitive functions, would be committed to the notion that differences in cognitive domains would be reflected primarily in differences in which brain regions support tasks in the domain. In contrast, redeployment-based approaches, being based on the idea that most brain regions are used in many different tasks across cognitive domains, would expect very little difference in which brain regions were used in each domain. However, because redeployment nevertheless expects brain regions to have fixed low-level functions [3–5], it is committed to the notion that differences in functions and domains must instead be the result of differences in the ways in which the areas cooperate in supporting different tasks. To put this in more concrete visual terms, imagine a simplified brain with six regions that together support two different cognitive domains. If one supports a localization-based (or a classical modular) organization for the brain, one would expect the regional cooperation patterns to look like those in the diagram on the left. In contrast, redeployment predicts an organization that looks something more like that shown in the diagram on the right (Fig. 2.2).




Fig. 2.2: Two different possibilities for the functional organization of the cortex. Figure shows an imagined brain with six regions supporting two cognitive domains. Localization predicts that domain 1 (*blue*) and domain 2 (*black*) will utilize different brain areas, while redeployment predicts that the domains will utilize many of the same brain areas, cooperating in different patterns.

There is an obvious analog for these features in our co-activation graphs: comparing the graphs from different domains, node overlaps indicate Brodmann areas that support tasks in both domains, whereas edge overlaps would indicate a similar pattern of cooperation between Brodmann areas. Thus, localization predicts little node

overlap between co-activation graphs (and therefore also low edge overlap), while redeployment predicts a great deal of node overlap, but little edge overlap. Using our database of imaging data, we did a co-activation analysis for the eight cognitive domains having more than 30 experiments: action; attention; emotion; language; memory; mental imagery; reasoning; and visual perception. The number of experiments (472 total) was not balanced between domains and authors, but otherwise followed the procedures outlined above. Using Dice's coefficient as our measure ($d = 2(o_{1,2})/(n_1 + n_2)$, where o is the number of overlapping elements and n is the total number of elements in each set), we compared the amount of node and edge overlap between each of the eight domains. As predicted by redeployment, we found a high degree of node overlap ($d = 0.81$, $SD = 0.04$) but very little edge overlap ($d = 0.15$, $SD = 0.04$). The difference is significant (two-sample Student's- t test, double-sided $p << 0.001$). Figure 2.3 shows a graph of the results. This is just one among a number of findings that suggest that redeployment is the better supported approach to understanding the functional topography of the cortex [4–6].




Fig. 2.3: Mean overlap of nodes vs. edges. A graph of the average Dice's coefficient for similarity between the sets of nodes and edges in a pair-wise comparison of co-activation graphs from eight cognitive domains. Difference between the means is significant ($p << 0.001$).

Looking at node and edge overlaps is just a simple example of the sorts of comparisons one might make using data in this format. Others more specific to graph-based representations also readily suggest themselves. For instance, one common form of analysis in graphs is a clique analysis, so called because of its origin in the analysis of social networks [2]. A clique is a maximal complete sub-graph – that

is, a set of nodes in a graph that are fully connected with one another, but not fully connected with any other node in the whole graph. In this context, a clique would indicate a set of Brodmann areas that are fully co-active with each other, but not with other areas of the brain; any such neural cliques would obviously be structures of interest. As in the case of social networks, however, this definition may be too strict for many purposes. Intuitively, we would be interested in sets of nodes that are cohesive and relatively isolated – that is, nodes that are highly but not necessarily fully connected, and much more connected with each other than with other nodes in the graph. These would represent sets of brain regions that are generally co-active with each other, but that operate with relative independence from the rest of the brain. Alba [2] offers the notion of a sociometric clique (an n -clique of diameter n), as well as measures of cohesiveness and isolation, that could be adopted here to discover sets of brain regions with the desired properties. Cohesive, isolated sociometric cliques seem likely to correspond to the neural components that cooperate to support a set of closely related cognitive functions or sub-functions. Whether this is so is an open scientific question, but such cliques are a far more plausible target for investigations into the neural components supporting particular cognitive functions than are individual brain areas. To return us to the issue with which this chapter began: co-activation graphs allow one to discover (among other things) neural cliques; in our view, what Anderson et al. should be doing is trying to match ACT-R modules to these sorts of structures, and not to individual brain areas.

These are far from the only research avenues that these data offer. One can also look at other features of the graphs, such as local topography, which may help make plausible inferences about underlying function. For instance, a hub-and-spoke pattern of co-activation may indicate broadcast or information consolidation functions; in contrast, long strings of connected nodes might indicate serial processing.

We could go on indefinitely, but the point is not to exhaustively list all the possible analyses one might make with graph-based co-activation data. Instead we would like to take the opportunity to call to mind the fact that, at very many points in the history of science, great progress has been made just in virtue of finding the right format for otherwise well-known data. In a field as young as Cognitive Neuroscience it is still more than possible for simple ideas to make a transformative impact; co-activation graphs may be one of those ideas.

2.5 Relating fMRI to EEG

We would like to conclude by describing one longer term application of co-activation graphs about which we are especially excited. As the reader is no doubt aware, a long-standing issue in experimental and clinical neuroscience has been the question of how to relate data from EEG/MEG to fMRI. Chief among the many obstacles standing in the way of relating the two have been (1) questions over whether each technology measures the same underlying neural activity [26] and (2) difficulty in finding the right representational format for the relation, given the vastly differ-

ent temporal scale of the two data streams [20]. However, recent research seems to indicate a mitigation of the first issue; and co-activation graphs may contribute to a novel approach to the second. We will discuss each of these in turn.

Although there have been for some time, and continue to be, questions about the neurophysiological bases of the fMRI signal, converging evidence strongly suggests that the BOLD signal is best correlated with local field potentials [25, 7, 35]. This is good news for the project of relating EEG and fMRI, because recent work has shown that EEG signals can also be analyzed to give estimations of LFP [29, 30]. Although this is hardly to be considered the last word on the subject, it appears that differences in underlying neurophysiological basis do not necessarily pose an obstacle to relating the two sources of data.

This brings us to the vast differences in temporal resolution. Since existing fMRI data cannot be made faster, typical solutions to the mismatch in temporal resolution have involved lowering the resolution of the EEG signal, by sampling signals over much longer timescales, and applying mathematical or statistical procedures (e.g., temporal averaging) to generate a relevant structure such as a local maximum in the 3D current distribution; this can then be compared to equivalent structures from fMRI. Vitacco et al. [36] applied this method to relate EEG and fMRI in a word classification task, but while they were able to obtain agreement between local maxima for group mean data, there was much poorer correspondence for individual subjects. One reason for this problem may be that, in averaging or otherwise manipulating EEG signals, one may be generating artifacts rather than discovering real features of the data. This is not to say that such attempts at data fusion are not promising, only that there is room for the introduction and evaluation of alternate approaches.

We have already outlined our approach to mining large numbers of fMRI studies and representing the results in graph format. This is relevant to the current issue because Chaovalitwongse et al. [13] recently developed a way to represent EEG data that also emphasized cooperative activity and also involved a graph-based representation scheme. In the scheme developed by Chaovalitwongse et al., cooperation between brain areas is measured in terms of the co-variance between EEG electrodes. Although the discovery of temporal correlation in large data sets is far from a trivial problem. Chaovalitwongse et al. [14, 12] have developed different methods to make such data mining tractable.

In discussions with Prof. Chaovalitwongse, we quickly realized that combining our two approaches could help address the issue of relating fMRI and EEG, because in approaches that focus on the cooperation of brain areas the small-scale temporal features of the EEG signal are de-emphasized, and the graph-based representational formats are entirely compatible; given the same underlying spatial segmentation of the cortex, the two cooperation graphs can be directly overlaid.

Of course, while it is clear that co-activation and co-variation graphs can be easily overlaid, what is unknown is whether there is any systematic relation between EEG co-variance and fMRI co-activation. We are currently putting together a research project to help answer this question (insofar as each graph is providing genuine information about which brain areas cooperate in supporting various cognitive tasks,

it certainly seems plausible that there would be some such relation). While it is by no means certain that any such relation will be found, the potential payoff is enormous. Among other things, it suggests it would be possible to mine the vast trove of fMRI data to provide baseline expectations for normal brain function in terms of the temporal correlation between brain areas. Since this can be observed cheaply, noninvasively, and in real time with EEG, it would be of great use in clinical settings for detecting deviations from normal function, such as might be observed prior to the onset of an epileptic seizure [12].

2.6 Conclusion

This chapter introduced a very simple analytical method for mining large numbers of brain imaging experiments to discover functional cooperation between brain regions. We reported some preliminary results of its application, illustrated some of the many future projects in which we expect the technique will be of considerable use, and described a research resource for investigating functional cooperation in the cortex that will be made publicly available through the lab web site. We hope and expect the availability of this resource will help spur new and innovative discoveries in the cognitive and computational neurosciences.

References

1. Abello, J., Pardalos, P.M., Resende, M.G.C. On maximum clique problems in very large graphs in external memory algorithms. In: Abello, J., Vitter, J. (eds.) *AMS-DIMACS Series on Discrete Mathematics and Theoretical Computer Science*, Vol. 50 (1999)
2. Alba, R.D. A graph-theoretic definition of a sociometric clique. *J Math Sociol* **3**, 113–126 (1973)
3. Anderson, J.R., Qin, Y., Jung, K.J., Carter, C.S. Information processing modules and their relative domain specificity. *Cogn Psychol* **54**, 185–217 (2007)
4. Anderson, M.L. Evolution of cognitive function via redeployment of brain areas. *Neuroscientist* **13**(1), 13–21 (2007)
5. Anderson, M.L. Massive redeployment, exaptation, and the functional integration of cognitive operations. *Synthese* **159**(3), 329–345 (2007)
6. Anderson, M.L. The massive redeployment hypothesis and the functional topography of the brain. *Philos Psychol* **21**(2), 143–174 (2007)
7. Attwell, D., Iadecola, C. The neural basis of functional brain imaging signals. *Trends Neurosci* **25**(12), 621–25 (2002)
8. Bock, R.D., Husain, S.Z. An adaptation of Holzinger's b-coefficients for the analysis of sociometric data. *Sociometry* **13**, 146–53 (1950)
9. Bonacich, P. Factoring and weighting approaches to status scores and clique identification. *J Math Sociol* **2**, 113–20 (1972)
10. Brannen, J.H., Badie, B., Moritz, C.H., Quigley, M., Meyerand, M.E., Haughton, V.M. Reliability of functional MR imaging with word-generation tasks for mapping Broca's area. *Am J Neuroradiol* **22**, 1711–1718 (2001)
11. Cabeza, R., Nyberg, L. Imaging cognition II: An empirical review of 275 PET and fMRI studies. *J Cogn Neurosci* **12**, 1–47 (2000)

12. Chaovalitwongse, W., Fan, Y.J., Sachdeo, R. On the k-nearest dynamic time warping neighbor for abnormal brain activity classification. *IEEE Trans Syst Man Cybern A Syst Hum* **37**(6), 1005–1016 (2007). To appear
13. Chaovalitwongse, W., Iasemidis, L.D., Pardalos, P.M., Carney, P.R., Shiau, D.S., Sackellares, J.C. Performance of a seizure warning algorithm based on the dynamics of intracranial EEG. *Epilepsy Res* **64**, 93–133 (2005)
14. Chaovalitwongse, W., Pardalos, P.M., Prokopyev, O.A. Electroencephalogram (EEG) time series classification: Applications in epilepsy. *Ann Operations Res* **148**, 227–250 (2006)
15. Diestel, R. *Graph Theory*, 3rd edn. Springer-Verlag, Heidelberg (2005)
16. Gross, J.L., Yellen, J. *Graph Theory and its Applications*, 2nd edn. Discrete Mathematics and Its Applications. Chapman & Hall/CRC, London (2005)
17. Han X., Xu, C., Braga-Neto, U., Prince, J.L. Topology correction in brain cortex segmentation using a multiscale, graph-based approach. *IEEE Trans Med Imaging* **21**, 109–121 (2002)
18. Hayes, B. Graph theory in practice: Part I. *Am Sci* **88**(1), 9–13 (2000)
19. Hayes, B. Graph theory in practice: Part II. *Am Sci* **88**(2), 104–109 (2000)
20. Horwitz, B., Poeppel, D. How can EEG/MEG and fMRI/PET data be combined? *Hum. Brain Mapp.* **17**, 1–3 (2002)
21. Laird, A.R., Fox, M., Prince, C.J., Glahn, D.C., Uecker, A.M., Lancaster, J.L., Turkeltaub, P.E., Kochunov, P., Fox, P.T. Ale meta-analysis: Controlling the false discovery rate and performing statistical contrasts. *Hum Brain Mapp* **25**, 155–164 (2005)
22. Laird, A.R., Lancaster, J.L., Fox, P.T. Brainmap: The social evolution of a functional neuroimaging database. *Neuroinformatics* **3**, 65–78 (2005)
23. Lancaster, J., Laird, A., Fox, M., Glahn, D., Fox, P. Automated analysis of meta-analysis networks. *Hum Brain Mapp* **25**, 174–184 (2005)
24. Lancaster, J.L., Woldorff, M.G., Parsons, L.M., Liotti, M., Freitas, C.S., Rainey, L., Kochunov, P.V., Nickerson, D., Mikiten, S.A., Fox, P.T. Automated talairach atlas labels for functional brain mapping. *Hum Brain Mapp* **10**, 120–131 (2000)
25. Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A. Neurophysiological investigation of the basis of the fMRI signal. *Nature* **412**, 150–157 (2001)
26. Nunez, P.L., Silberstein, R.B. On the relationship of synaptic activity to macroscopic measurements: Does co-registration of EEG with fMRI make sense? *Brain Topogr* **13**, 79–96 (2000)
27. Sporns, O., Ktter, R. Motifs in brain networks. *PLoS Biol* **2**, e369 (2004)
28. Özcan, M., Baumgärtner, U., Vucurevic G. Stoeter, P., Treede, R.D. Spatial resolution of fMRI in the human parasyllvian cortex: Comparison of somatosensory and auditory activation. *NeuroImage* **25**(3), 877–887 (2005)
29. Grave de Peralta Menendez, R., Gonzales Andino, S., Morand, S., Michel, C., Landis, T. Imaging the electrical activity of the brain. *Electra Hum Brain Mapp* **9**, 1–12 (2000)
30. Grave de Peralta Menendez, R., Murray, M.M., Michel, C., Martuzzi, R., Gonzales Andino, S.L. Electrical neuroimaging based on biophysical constraints. *NeuroImage* **21**, 527–539 (2004)
31. Sporns, O., Tononi, G., Edelman, G.M. Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cereb Cortex* **10**, 127–141 (2000)
32. Suharitdamrong, W., Chaovalitwongse, A., Pardalos, P.M. Graph theory-based data mining techniques to study similarity of epileptic brain network. In: *Proceedings of DIMACS Workshop on Data Mining, Systems Analysis, and Optimization in Neuroscience* (2006)
33. Turkeltaub, P.E., Eden, G.F., Jones, K.M., Zeffiro, T.A. Meta-analysis of the functional neuroanatomy of single-word reading: Method and validation. *Neuroimage* **16**, 765–780 (2002)
34. Ugurbil, K., Toth, L., Kim, D.S. How accurate is magnetic resonance imaging of brain function? *Trends Neurosci.* **26**(2), 108–114 (2003)
35. Viswanathan, A., Freeman, R.D. Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nat Neurosci* **10**(10), 1308–1312 (2007)
36. Vitacco, D., Brandeis, D., Pasual-Marqui, R., Martin, E. Correspondence of event-related potential tomography and functional magnetic resonance imaging during language processing. *Hum Brain Mapp* **17**, 4–12 (2002)