ACADEMIC PRESS

Available online at www.sciencedirect.com
SCIENCE DIRECT®

NeuroImage 19 (2003) 457–465

NeuroImage

www.elsevier.com/locate/ynimg

A report of the functional connectivity workshop, Dusseldorf 2002

Lucy Lee,* Lee M. Harrison, and Andrea Mechelli

Wellcome Department of Imaging Neuroscience, Institute of Neurology, London WC1N 3BG, UK

Received 5 September 2002; accepted 10 December 2002

Abstract

This report provides a commentary on the issues presented and discussed at the recent "Functional Brain Connectivity" workshop, held in Dusseldorf, Germany. The workshop brought together researchers using different approaches to study connectivity in the brain, providing them with an opportunity to share conceptual, mathematical, and experimental ideas and to develop strategies and collaborations for future work on functional integration. The main themes that emerged included: (1) the importance of anatomical knowledge in understanding functional interactions the brain; (2) the need to establish common definitions for terms used across disciplines; (3) the need to develop a satisfactory framework for inferring causality from functional imaging and EEG/MEG data; (4) the importance of analytic tools that capture the dynamics of neural interactions; and (5) the role of experimental paradigms that exploit the functional imaging of perturbations to cortical interactions.

1. Introduction

The "Functional Brain Connectivity" workshop held recently in Dusseldorf, Germany, brought together prominent researchers using different approaches to study connectivity and allowed them to exchange conceptual, mathematical, and experimental ideas. It was hoped that by inviting researchers with diverse areas of interest, new insights would be gained into the study of functional integration within the human brain. To encourage discussion, the speakers were asked to limit the formal section of their presentation to 15 minutes and to use just two transparencies or slides; consequently much of the important information emerged during the discussions. It is hoped that this report provides a fair representation of the views expressed by participants during the presentations and discussions.

The report is divided into three sections: (1) connectivity: definitions, causality, and inference; (2) development of experimental techniques; and (3) Analytic techniques and modeling.

2. Connectivity: definitions, causality, and inference

Three different types of connectivity were discussed during the workshop: anatomical, functional, and effective connectivity. This section focuses on the definitions and limitations of these concepts and their relevance to the study of how the brain works, particularly with respect to the current shift in emphasis from studies of functional segregation to functional integration.

2.1. Anatomical connectivity

Although the definition of the term "anatomical connectivity" was not directly addressed, discussions about anatomical connections encompassed two main themes: first, the role of anatomical data in modelling neuroimaging data and constraining computational simulations, and second the boundary between anatomical description and functional characterisation of connections. The first of these two issues was addressed directly by the provocative question: Why do we need to know any anatomy? Would it be acceptable to simply infer the presence of an anatomical connection from the functional characteristics of a system? The consensus reached was that knowledge of anatomy is important to define the "connectivity space" thereby providing plausible

* Corresponding author. Fax: +44-0-207-813-1420.
E-mail address: llee@fil.ion.ucl.ac.uk (L. Lee).

1053-8119/03/$ – see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S1053-8119(03)00062-4

biological constraints for theories and inferences about neural interactions when analysing functional neuroimaging data and developing computer simulations. Having agreed that some knowledge of anatomy is important for all levels of inquiry, a new question was raised: What measurements of the anatomical connectivity are most useful to the study of how the brain works? Knowing if there are direct connections between two neurones or cortical areas is clearly important, but a complete description of the connections includes information such as the receptor subtypes at synapses (e.g., AMPA vs. NMDA), the ratio of inhibitory to excitatory interneuron targets, the number of connections, and their physiological impact (modulatory top-down vs. driving bottom-up inputs). At this point the boundary between anatomy and function becomes blurred. For example, Peter Buzas suggested that the cells he investigates in primary visual cortex might be better classified if they were divided according to their responses to physiological stimuli rather than their morphology, that is, functional instead of anatomical criteria (see Section 3.3). Karl Zilles further emphasised this point, reminding participants that labelling of macaque cortical areas, for example, V2-5, is based on functional characteristics and it is almost impossible to distinguish among these areas using anatomical criteria alone.

2.2. Functional segregation versus functional integration

Characterising brain activity in terms of the functional specialisation of sensorimotor and higher cognitive brain areas is the primary approach to functional neuroimaging data. However, characterising brain activity in terms of functional specialisation does not reveal anything about how different brain regions communicate with each other, despite the fact that communication is an implicit assumption of this specialisation. For instance, the neural correlates of semantic processing can be identified using written words by virtue of the assumption that visual regions interact with semantic regions. Thus, analyses of neural activity based solely on functional specialisation provide only a limited account of the neuronal substrate of the processes investigated. Alternative approaches have therefore been developed, in the context of functional neuroimaging, to investigate the integration of functionally specialised areas (see Section 4.1).

2.3. Functional and effective connectivity

The approaches mentioned in Section 2.2 use a number of concepts and definitions derived from multiunit microelectrode recording of separable spike trains (Aertsen and Preissl, 1991; Gerstein and Perkel, 1969). In this context, *functional connectivity* is defined as the "temporal correlations between spatially remote neurophysiological events" (Friston et al., 1993a), whereas *effective connectivity* is defined as "the influence that one neural system exerts over another either directly or indirectly" (Friston et al., 1993b). At the level of multiunit microelectrode recordings, functional connectivity can result from stimulus-locked transients, evoked by a common afferent input, or may reflect stimulus-induced phasic coupling of neuronal assemblies, mediated by synaptic connections among brain areas. Ed Bullmore pointed out that there are practical differences in the characterisation of functional and effective connectivity in multiunit electrode recording and functional neuroimaging techniques, specifically issues of spatial and temporal resolution, and that this may have implications for the relevance of these definitions across disciplines.

During the initial discussions Karl Friston expanded on the motivation behind these definitions of functional and effective connectivity within the context of functional neuroimaging: namely, to emphasise the difference between *descriptions of patterns* of neural activity and possible *explanations of their origins.* Functional connectivity reduces to testing the null hypothesis that activity in two regions shares no mutual information. Mutual information is a statistical description of the degree to which two regions demonstrate similar behaviour or statistical interdependence (Cover and Thomas, 1991). In other words, the characterisation of brain activity in terms of functional connectivity is "model free." In contrast, characterising brain activity in terms of effective connectivity requires a causal or acausal model, in which regions and connections of interest are specified by the researcher, often constrained by a combination of neuroanatomical, neuropsychological, and functional neuroimaging data. This is a crucial point when considering the distinction between functional and effective connectivity because it emphasises the shift from a description of what the brain does to a theory of how it does it.

However, there has not always been precise and universal agreement on the definitions and concepts. For example, early PET studies exploited intersubject variability to make inferences about functional connectivity (see McIntosh and Gonzalez-Lima, 1994; Horwitz et al., 1998). Here it was assumed that differences across subjects lead to correlated activities that are an index of functional integration. Clearly, this measure of connectivity discounts temporal correlations and does not conform to the definitions given above. In contrast, most fMRI studies, like EEG and MEG, have focussed on temporal correlations by treating data from different subjects as if they came from the same subject (see Iidaka et al., 2001; Bokde et al., 2001; Bullmore et al., 2000; Maguire et al., 2000; Coull et al., 1999). However, estimates of the coupling can still be influenced by item-to-item or block-to-block variability over subjects using this approach. In short, there is sometimes a tension between the definition of what is being measured and how it is measured.

2.4. Problems associated with inferring causality from functional neuroimaging data

It is difficult to distinguish unambiguously between functional and effective connectivity simply by looking at the

data because temporal correlations between neurophysiological events in separate neural systems (functional connectivity) may or may not be due to the influence that one neural system exerts over another (effective connectivity). Ed Bullmore discussed two possible ways of inferring effective connectivity (temporal precedence and perturbation studies) and highlighted the problems associated with each of them.

2.4.1. Temporal precedence

One method of inferring causality involves the principle of "temporal precedence": if activity in area A occurs prior to activity in area B, then activity in A might cause activity in B through connections between the two areas. However, in the context of functional MRI (fMRI), temporal precedence at the neuronal level may be masked at the haemodynamic level, because of the temporal smoothing inherent in the coupling between synaptic activity and haemodynamic changes. In extreme cases, one region may cause activity in another and yet, because of different latencies, the haemodynamic response in the "target" region may be expressed *before* the haemodynamic response in the "source" region. Even EEG and MEG studies, which provide a temporal resolution in the order of milliseconds, may not provide a satisfactory solution since temporal precedence is not unambiguously linked to causality. Despite this note of caution, it was agreed that EEG and MEG studies are important in the study of functional integration and causality, and this was reflected by the number of presentations describing the application of time-series analysis methods to these data (see Section 4.3).

2.4.2. Perturbing the system

Perturbing the system might provide a second way of inferring causality. The majority of functional neuroimaging experiments involve experimental manipulations of neural activation (i.e., perturbations), in the form of stimuli, be they visual, auditory, psychological, etc. There are two important advantages of using perturbation studies. Firstly, the manipulation is under explicit experimental control and can therefore be located precisely in time and space. Secondly, it is possible to selectively alter neural activity in specific cortical areas and to assess the effect that this has on the activity and interactions between the nonperturbed areas. A variety of perturbation studies where direct electrical stimulation of exposed cortex, cortical cooling, and transcranial magnetic stimulation (TMS) are combined with different functional imaging techniques are discussed in Section 2.4.

2.5. Conclusions and future research

Although functional imaging experiments are generally based on the principle of functional specialisation, this type of analysis does not provide information about how different brain regions communicate with each other. In contrast, the analysis of functional and effective connectivity addresses neuronal interactions. Although functional connectivity and effective connectivity can be invoked at a conceptual level in both electrophysiology and neuroimaging, the nature and the scale of the neurophysiological measures provided by these techniques differ considerably. The combination of simultaneously acquired "microscopic" measures (obtained using multiunit electrode recording) and "macroscopic" measures (obtained through PET and fMRI scanning) is desirable for a better understanding of functional integration in the brain. There was general agreement that this has not yet been achieved and remains an important area for further research.

3. Development of experimental techniques

The informal atmosphere of this workshop provided an opportunity for researchers to outline new areas of research and present preliminary results. In this section many of the technical developments are introduced with a superficial level of description, with emphasis placed on their potential contribution to the study of functional integration. The section is divided according to four major themes that emerged during the discussions, many of which span a range of methods, from characterisation of microscopic corticocortical connections to perturbation studies in human subjects.

3.1. Bridging the gap between different species

One of the themes discussed on several occasions was the "information gap" that exists between the wealth of anatomical and single-unit recording data in monkeys and the paucity of information about the micro- and macroscopic connections within the human central nervous system. Two presentations in particular highlighted different approaches to this problem.

The approach taken by Wim Vanduffel has been to perform fMRI studies at 1.5 T in awake, behaving monkeys (Vanduffel et al., 2001). In this way, data from nearly identical experimental paradigms (visual motion processing for example) can be compared between monkeys and humans, as well as enabling fMRI data to be compared to single-unit work within the same species.

Karl Zilles presented results from experimental techniques developed to study anatomical connectivity of postmortem human brains, including differential myelin staining of the optic radiation (Burgel et al., 1997) and autoradiographic imaging of receptor subtypes within the cortex and basal ganglia (Morosan et al., 2001; Rademacher et al., 2001; Geyer et al., 2001). Nonlinear warping (Schormann and Zilles, 1998) is being used to compile this anatomical data from the human brain into a multimodal 3-D representation with a single spatial reference system, from microscopic, cellular detail to the surface anatomy and white matter tracts (Mazziotta et al., 2001; Roland et al.,

2001). Using the same spatial reference system, data from functional imaging studies can be interpreted with respect to probabilistic population anatomy (Binkofski et al., 2002; Bodegard et al., 2001).

3.2. Diffusion tensor imaging

The relatively new field of diffusion tensor imaging (DTI) was the subject of several presentations and discussions. The consensus reached was that DTI is an exciting and potentially valuable technique, even if it is unlikely to reach the resolution of single neurones or synaptic connections and therefore will not replace techniques such as fibre tracing. Indeed, Ed Bullmore suggested that the population maps of human white matter tract distributions described above might be used as priors for DTI tracking algorithms.

Several presentations discussed validation of DTI results: the awake monkey fMRI setup of Wim Vanduffel enables DTI results in monkeys to be compared to current anatomical knowledge of the macaque. Rainer Goebel showed DTI data from anaesthetised cats, which have been used to compare the performance of a variety of different tracking algorithms, the results of which will be compared to postmortem tract tracing experiments in the same animals (Kim et al., 2001). Goebel's work uses fMRI to locate cortical areas activated during visual stimulation which are then used as starting points for tracking. In this approach the functional data are used to constrain the estimation of anatomical connections. Martin Koch presented a similar experiment using DTI and a Monte Carlo simulation to estimate the probability of two functionally correlated voxels, within a single slice of human cortex, being anatomically connected (Koch et al., 2002).

The utility of DTI as a research tool will depend on how reliable and reproducible these various tracking methods prove to be. The potential applications for information from DTI are currently unexplored. Firstly, DTI may be clinically useful for assessing subtle changes in white matter anatomy, for example, during recovery from brain lesions. Secondly, in analyses of effective connectivity DTI data may be useful as a constraint on possible anatomical connectivity, but it is unlikely to provide information about the direction or functional classification of connections. It should be noted that not all participants thought DTI was a useful technique, proponents of DTI were reminded on several occasions that DTI did not measure "connections" but simply large tracts of similarly orientated myelinated axonal processes. This scepticism about DTI reflected the deeper question: Does macroscopic anatomy add anything to a complete characterisation of effective connectivity in terms of physiology and function?

3.3. Using functional characteristics to infer anatomical connections

The use of functional data to make inferences about the presence of anatomical connections is not confined to macroscopic brain mapping. Dirk Schubert and Rolf Kotter presented new results, extending their previously published work exploring the relationship between the functional characteristics of cells in rat barrel cortex and the synaptic connections between different cell types within cortical layers of the same barrel and between neighbouring barrels (Schubert et al., 2001). Several interesting suggestions were made during the discussion of their work, including the idea of extending the use of caged glutamate to include specific agonists and antagonists and the use of caged GABA.

Peter Buzas presented work investigating the integration of functional characteristics and lateral cortical connections within primary visual cortex by combining anatomical data from three-dimensional reconstruction of cortical cells with maps of orientation, direction, and ocular dominance preference, obtained by optical imaging (Buzas et al., 1998, 2001). His results revealed profound differences in the functional characteristics of the lateral connections between cells with similar morphology and anatomical connectivity, raising the issue of how best to classify these cells: according to morphological or functional features.

Mark Hubener presented preliminary results from experiments using brief pulses of electrical stimulation to the cat primary visual cortex to generate patterns of increased neural activity that were subsequently compared to responses to visual stimuli (gratings) with different spatial orientations using optical imaging. The striking similarities in size and location of the activated areas suggests that the functional maps, generated in response to physiological stimuli, depend in part upon the same anatomical connections that are responsible for the patterns seen after direct electrical stimulation. Tomas Paus suggested during the discussion that if it were possible to severe the horizontal corticocortical connections surrounding the stimulating electrode this would provide some measure of their contribution to the effects of electrical stimulation.

3.4. Perturbation studies

Experimental techniques that involve destruction of anatomical connections or experimental manipulation of the excitability of cortical areas formed another important theme. The ability to experimentally activate or deactivate an area of the brain during a task, and during functional imaging, is a powerful method for exploring effective connectivity and causality. Wim Vanduffel indicated they are working towards incorporating reversible lesion studies such as cortical cooling with their monkey fMRI experiments.

Tomas Paus discussed transcranial magnetic stimulation (TMS) as a technique for altering neural activity in humans under explicit experimental control. He described a series of experiments combining TMS with PET where changes in neural activity secondary to artificial stimulation where compared with changes evoked by physiological stimuli (Paus et al., 1997, 1998). The ad-

vantages of combining TMS with functional imaging were highlighted: first, as TMS is independent of behaviour, observed changes in neural activity are not confounded by the subjects' ability to perform a task or by the strategy used. Second and perhaps of more interest, the combination of TMS with PET can be used to alter the activity of one brain area and observe the effects on activity in other areas, either in response to further TMS or during behavioural paradigms. This experimental strategy lends itself very readily to an analysis of effective connectivity (see Section 2.4).

3.5. Comments

The division of this section according to broad themes, not into micro- and macroscopic techniques runs the risk of creating artificial divisions between various lines of inquiry, and grouping together other techniques that do not necessarily reflect similar underlying scientific questions. For example, although DTI and tract tracing/differential myelin staining are very different techniques, and direct comparisons between them are controversial (see Section 3.2), they represent macroscopic and microscopic approaches to the same problem: establishing the presence of anatomical connections within the human brain. Despite differences in opinion between participants, there was general agreement that these techniques were complementary, and it was possible to envisage overlap between them in the future. In contrast, despite similarities in experimental design and statistical/inferential techniques (for example, the experiments described by Mark Hubener and Tomas Paus), the different methods of quantifying micro- and macroactivations mean that these experiments appear to address very different aspects of structure and functional integration. This apparent lack of overlap between microscopic and macroscopic approaches is reflected in the unresolved questions about the level of detail required to explain and model various brain mechanisms.

4. Analytic techniques and modeling

With the development of new techniques for measuring different aspects of brain activity comes an increasing demand for innovative methods of analysing data to finesse our understanding of how the brain works. The brain's connectivity, both anatomical and functional, is highly plastic, given the changing environment to which it adapts, for example, during development and learning or in response to disease. An important objective is to develop robust metrics and convincing explanations of the dynamic nature of neural connectivity associated with sensorimotor and cognitive processes in health and disease.

4.1. Models of effective connectivity

Randy McIntosh discussed the central issue of plasticity within large-scale neural networks. Early experiments using PET, and modeled with structural equation modelling (McIntosh and Gonzalez-Lima, 1994), demonstrated stimulus-dependent changes of interregional connectivity, exploiting time-independent features of the data. These experiments were crucial in establishing large-scale networks of cortical activity as a possible substrate of cognitive or sensorimotor tasks. They also provided a simple measure of stimulus-dependent short-term plasticity using functional neuroimaging techniques. Coordinated cortical networks, subserving behaviour, provide a source of diversity for neural representations, which has led to the notion of neural context (McIntosh, 1999). Cortical regions activate similarly within different tasks. For example, area 46 activated during visual object and spatial recognition is distinguished by activities within connected regions or the neural context within which the common region is influenced. A functionally specialised region then becomes dependent on its network of cortical influences and no longer has a functionally static response. Structural equation modelling initiated efforts to model such stimulus-dependent neural dynamics, with the development of increasingly sophisticated models being the topic of further discussions during the workshop.

Karl Friston outlined a synthesis of models developed over the past decade and demonstrated how they were special cases of a general dynamic description of the brain. He made the distinction between the use of statistics to describe correlations among neural activations and modeling causal structures that could explain these observations. Treating the brain as a dynamic system from which we are able to observe certain physiological parameters (changes in BOLD signal, electrical or magnetic signals) over time, given particular inputs (experimentally designed perturbations of the system), we can attempt to estimate parameters that govern the dynamics of underlying states, on which these observations are dependent. Starting from the general state and observation equations describing the measured temporal evolution of a system, expressions of effective connectivity were derived, demonstrating that structural equation modelling (Buechel and Friston, 1997), multivariate autoregressive models, and the Volterra approach (Friston, 2001) were all special cases of dynamic causal modelling.

The talk ended with a brief description of dynamic causal modeling (DCM), describing coupling between interacting states within the general framework of multiple input, multiple output (MIMO) systems. In its current form the state equations are expanded according to a bilinear approximation, with the linear terms describing "latent" couplings among the states in the absence of input, and bilinear terms estimating the modulation of these couplings by inputs to the system. The procedure reduces the highly nonlinear dynamics of the brain to a tractable linear problem of

estimating couplings given experimental causes (inputs) and measured responses. Bayesian methods are used to estimate coupling parameters contained within the state equations. During the discussion of this work, clarifications were made regarding the scope of these analyses: firstly, that they are useful for designed experiments where the investigator is only interested in outputs that vary with the experimental manipulations and, secondly, that these methods are not applicable to systems with noncontrollable dynamic elements, that is, systems with their own intrinsic activity, such as epileptic foci.

4.2. The use of simulation modelling to study effective connectivity

The contribution from Barry Horwitz raised a fundamental issue about the biological plausibility of the models used for effective connectivity. Horwitz reviewed a series of realistic and carefully constructed forward models, based on neurophysiological principles that could be used to generate simulated imaging data. Furthermore, such models were able to accommodate experimental perturbations of the sort discussed by Thomas Paus, namely TMS, as well as standard cognitive and sensorimotor challenges. The critical utility of this approach is that it explicitly models the underlying neuronal substrate of measured responses and vicarious measures of effective connectivity. More generally, the approach described by Horwitz afforded us tenable forward models of observed imaging data that embodied realistic neuronal dynamics and interregional integration. A good example of this approach can be found in Husain et al. (2002).

4.3. Time-series analysis

A recurrent theme throughout the presentations was the use of time-series analysis tools to describe and model the brain as a dynamic system. Time-series analysis techniques represent a possible way of extracting dynamic characterisations of the underlying state variables responsible for our observations (BOLD and electrical/magnetic signals). Some of these methods are currently used to analyse EEG/MEG and LFP recordings (Bressler et al., 1999) and are becoming increasingly popular in fMRI and PET data analysis.

Jim Stone described independent component analysis (ICA), a blind source separation technique (BSS) (Stone, 2002, and www.shef.ac.uk/~pcljvs, which includes demonstrations and Matlab code). The problem faced by experimentalists measuring signals produced from different sources within the brain is that data collected at a distance, for example, at the scalp with EEG, are mixtures of signals from different sources. The challenge is then to separate a single source from the rest, the classic analogy being a scenario from a cocktail party where one can imagine isolating a single voice in contrast to the others that merge into one indistinguishable background noise. ICA is a method for unmixing measured time series under the constraint that individual sources are statistically independent, the output of the algorithm being spatial or temporal modes of activity whose independence has been maximised. Information theoretic arguments concerning maximally efficient coding suggest the brain might use similar schemes, where a population of neurons encodes information with minimal redundancy, which is equivalent to maximising the independence between neurons.

Nicholas Schiff described a time-series analysis procedure called hierarchical decomposition (Repucci et al., 2001; see also www.users.med.cornell.edu/~jdvicto/nlardf.html#introduction) to characterise a “fingerprint” of the underlying non-linear dynamics that may be generating EEG signals during temporal lobe epilepsy (TLE) and absence seizures. Crucially, this approach aims to bridge the gap between observations of measured brain activity at the scalp surface and possible explanations of the hierarchical structure of electrical generators within the brain.

Schiff demonstrated the procedure on multichannel EEG recordings during a TLE seizure. The dimensionality of the data is reduced with principle component analysis and parameterised using a multivariate linear autoregressive model, before rotating the time series into a hierarchical structure. Each time series is then overfitted with a linear autoregressive (AR) model (model order 20) and what is left is modeled as a nonlinear AR process, an AR process being one whose current state is dependent on previous states in its recent history. A second-order AR bilinear term estimated from this procedure is then tested for goodness of fit and plotted on a two-dimensional graph, with each axis representing a different time lag, up to 20 lags. The result is what Schiff describes as a fingerprint of nonlinearities. The interesting finding from his work is that, despite gross differences in the raw EEG data from TLE and absence seizures, they share similar fingerprints, raising the question of whether they are generated by similar dynamics within the brain. If the method has succeeded in capturing a characteristic feature of underlying nonlinear coupling between oscillatory generators responsible for the globally synchronous activity during seizures, then could such characterisations exist for particular cognitive task? The participants agreed with Schiff that this is an exciting prospect.

4.4. Investigating the links between anatomy and function using computational simulations

Olaf Sporns presented his work on the relationship between structure and function of complex systems using computational simulations to explore the principles underlying their dependence (Sporns et al., 2000). The objective was to investigate the dialogue between structure and function through synthetic models of functionally realistic networks, looking for emergent structural “motifs” and similarities with real neural structures. Considering a system composed of individual units, interacting among themselves

to varying degrees, clusters of activity emerge. Such clusters can be considered as functionally specialised aggregates of activity reminiscent of neuroanatomical findings, with regional specialisation to specific sensorimotor or cognitive functions, for example, vision (Zeki, 1978). However, a balance exists between the amount of specialisation and the degree to which these regions functionally integrate. Sporns described a mathematical property called "complexity," which is a metric based on the mutual information among hierarchical bipartitions of a system. The metric is dependent on the balance between functional specialisation and integration, being zero at the extremes and maximal in between. Artificially generated "complex" systems reveal structural motifs that share salient features with real neural systems, in that functional clustering is a consistent feature of optimally complex systems. Such measures of real networks (using "wiring diagrams" of cat and monkey brains) demonstrate that these neural systems exist at a near maximal level of complexity and that this can be reduced through simulating disconnections.

4.5. A note on reductionism

As mentioned in Section 2.4, the issue of what we are willing to accept as a causal explanation of our observations was a persistent theme of the discussions and Ed Bullmore's presentation. Scientific methodology traditionally explores mechanistic explanations of causal structure. Phenomena are understood through a process of reducing the whole to its constituent parts and building up a mechanistic understanding of each step. However, given highly complex organisations of matter such as the brain this approach is insensitive to principles of collective order that emerge from self-organising systems. These are not explained by reducing the system to its individual parts as the very nature of emergent dynamics is an expression of the whole. Such a holistic perspective produces fundamentally different accounts of causal structure. Synergetics (Haken, 1983) was given as an example of such an approach.

Synergetics is a theory of self-organisation, a ubiquitous phenomenon throughout nature. The dynamics of "many-body" systems as they go through periods of instability (bifurcation) from one stable state to another (phase transition) is often governed by a few variables, called order parameters. This can be understood through the slaving principle. Near to a bifurcation the slowly varying order parameters enslave rapidly relaxing stable modes. Consequently the dynamics of the order parameters capture the collective behaviour of phase transitions without modelling microscopic properties explicitly. The benefit of this formulation is that the order parameter equation is typically low dimensional compared to the entire system.

Synergetic principles have been applied to EEG/MEG data analysis, describing long-range spatiotemporal patterns induced by cognitive and sensorimotor tasks (Kelso and Haken, 1997). Mean field approaches (e.g., Treves, 1993) apply a similar holistic approach to modelling macroscopic neural behaviours. In the same way that statistical thermodynamics characterises the macroscopic property of temperature in terms of distributions of an ensemble of microscopic units (atoms) within a system, such approaches aim to describe spatiotemporal neural patterns in terms of statistical distributions of microscopic states such as neuronal firing.

Accepting holistic explanations requires a different paradigm of what constitutes a sufficient framework compared to step-by-step, mechanistic accounts. No doubt both approaches will be important to our emerging notions of brain function.

5. Conclusions

This workshop recognised the significant progress made in recent years in understanding the anatomical structure and functional architecture of the central nervous system and provided insights into recently developed experimental and analytical methods for characterising functional integration. By bringing together researchers from diverse backgrounds it was possible to gain new insights into the strengths and weaknesses of their different approaches and to identify important directions for future research.

The important issues arising from the workshop are summarised below:

1. Anatomical detail has an important role in the study of functional integration within the brain, and this includes information about characteristics of these connections such as the types of receptors and neurotransmitters.
2. The distinction between functional and effective connectivity, as implied by the definitions framed by Karl Friston, is an important conceptual tool to select the appropriate methods for characterising functional integration.
3. There is a need to define a satisfactory framework for inferring causality at the temporal and spatial resolution of EEG, MEG, and fMRI data. This may be helped by the development of experimental paradigms to explicitly address the issue of effective connectivity and to reduce the divide between electrophysiology and functional imaging. These include further work on combined single-unit recording and fMRI and perturbation studies such as combined electrical stimulation with optical imaging, reversible deactivation in monkey fMRI studies, and the combination of TMS with EEG, PET, and hopefully fMRI in human subjects.
4. Statistical and modeling frameworks such as hierarchical decomposition and DCM that enable the nonlinear dynamics of neural interactions to enter analyses of brain activity are important developments and

may play a role in the study of functional integration and in addressing issues of causality.

Acknowledgments

The authors would like to thank Barry Horwitz and Rolf Kotter for their helpful comments and perspectives on this report.

Appendix

Organizers

Rolf Kotter (Heinrich Heine University, Dusseldorf, Germany)

Karl Friston (University College London, London, UK)

Sponsors

Boehringer Ingelheim Foundation

EU Thematic Network "Computational Neuroscience and Neuroinformatics"

Heinrich Heine University Düsseldorf

Universitätsklinikum Düsseldorf

Speakers

Ed Bullmore (University of Cambridge, Cambridge, UK): Significance of connectivity differences between groups of fMRI data.

Peter Buzas (Ruhr-Universitat, Bochum, Germany): Fine structure of cortical connections—fine structure of cortical maps.

Karl Friston (University College London, London UK): Dynamic causal modelling.

Rainer Goebel (University of Maastricht, Maastricht, Netherlands): Combining fMRI and DTI—constraints for functional connectivity analysis.

Barry Horwitz (National Institutes of Health, Bethesda, MD): Using simulation modeling to determine the neural correlates of functional connectivity.

Mark Hubener (Max-Planck-Institute of Neurobiology, Martinsried, Germany): Studying connectivity with optical imaging.

Martin A. Koch (University of Hamburg, Hamburg, Germany): An investigation of functional and anatomical connectivity using MRI.

Randy McIntosh (University of Toronto, Toronto, Ontario, Canada): Linking cognitive function and brain through neural context.

Tomas Paus (McGill University, Montreal, Quebec, Canada): Imaging and stimulating the human brain.

Nicholas Schiff (Cornell University, New York, NY): Quantitative approaches to system identification in human epilepsies.

Dirk Schubert and Rolf Kotter (Heinrich Heine University, Dusseldorf, Germany): Mapping functional connectivity through optical stimulation.

Olaf Sporns (Indiana University, Bloomington, IN): From neuroanatomy to functional brain connectivity—and back?

Jim V. Stone (Sheffield University, Sheffield, UK): Independent component analysis—principles and applications.

Wim Vanduffel (Leuven Medical School, Leuven, Belgium): Ground truth for functional connectivity measures.

Karl Zilles (Institute of Medicine, Research Centre Jülich, Jülich, Germany): Contribution of postmortem studies to functional and diffusion tensor imaging.

References

Aertsen, A., Preissl, H., 1991. Dynamics of Activity and Connectivity in Physiological Neuronal Networks. VCH, New York.

Binkofski, F., Fink, G.R., Geyer, S., Buccino, G., Gruber, O., Shah, N.J., Taylor, J.G., Seitz, R.J., Zilles, K., Freund, H.J., 2002. Neural activity in human primary motor cortex areas 4a and 4p is modulated differentially by attention to action. J. Neurophysiol. 88, 514–519.

Bodegard, A., Geyer, S., Grefkes, C., Zilles, K., Roland, P.E., 2001. Hierarchical processing of tactile shape in the human brain. Neuron 31, 317–328.

Bokde, A.L., Tagamets, T.A., Friedman, R.B., Horwitz, B., 2001. Functional interactions of the inferior frontal cortex during the processing of words and word-like stimuli. Neuron 30, 609–617.

Bressler, S.L., Ding, M., Yang, W., 1999. Investigation of co-operative cortical dynamics by multivariate autoregressive modeling of event-related local field potentials. Neurocomputing 26–27, 625–631.

Buechel, C., Friston, K.J., 1997. Characterising functional integration, in: Frackowiak, R.S.J., Friston, K.J., Frith, C.D., Dolan, R.J., Mazziotta, J.C., (Eds.), Human Brain Function, Academic Press, San Diego, pp. 127–140.

Bullmore, E., Horwitz, B., Honey, G., Brammer, M., Williams, S., Sharma, T., 2000. How good is good enough in path analysis of fMRI data? NeuroImage 11, 289–301.

Burgel, U., Mecklenburg, I., Blohm, U., Zilles, K., 1997. Histological visualization of long fiber tracts in the white matter of adult human brains. J. Hirnforsch. 38, 397–404.

Buzas, P., Eysel, U.T., Adorjan, P., Kisvarday, Z.T., 2001. Axonal topography of cortical basket cells in relation to orientation, direction and ocular dominance maps. J. Comp. Neurol. 437, 259–285.

Buzas, P., Eysel, U.T., Kisvarday, Z.F., 1998. Functional topography of single cortical cells: an intracellular approach combined with optical imaging. Brain Res. Brain Res. Protocols 3, 199–208.

Coull, J.T., Büchel, K.J., Friston, K.J., Frith, C.D., 1999. Noradrenergically mediated plasticity in a human attentional neuronal network. NeuroImage 10, 705–715.

Cover, T.M., Thomas, J.A., 1991. Elements of Information Theory. Wiley, New York.

Friston, K.J., 2001. Brain function, non-linear coupling, and neuronal transients. Neuroscientist 7, 406–418.

Friston, K.J., Frith, C.D., Frackowiak, R.S.J., 1993a. Time-dependent changes in effective connectivity measured with PET. Hum. Brain Mapp. 1, 69–80.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J., 1993b. Functional connectivity: the principal component analysis of large (PET) data sets. J. Cereb. Blood Flow Metab. 13, 5–14.

Gerstein, G.L., Perkel, D.H., 1969. Simultaneously recorded trains of action potentials: analysis and functional interpretation. Science 164, 828–830.

Geyer, S., Schleicher, A., Schormann, T., Mohlberg, H., Bodegard, A., Roland, P.E., Zilles, K., 2001. Integration of microstructural and functional aspects of human somatosensory areas 3a, 3b, and 1 on the basis of a computerized brain atlas. Anat. Embryol. (Berlin) 204, 351–366.

Haken, H., 1983. Synergetics: An Introduction, third ed. Springer-Verlag, New York.

Horwitz, B., Rumsey, J.M., Donohue, B.C., 1998. Functional connectivity of the angular gyrus in normal reading and dyslexia. Proc. Natl. Acad. Sci. USA 95, 8939–8944.

Husain, F.T., Nandipati, G., Braun, A.R., Cohen, L.G., Tagamets, M.A., Horwitz, B., 2002. Simulating transcranial magnetic stimulation during PET with a large-scale neural network model of the prefrontal cortex and the visual system. NeuroImage 15, 58–73.

Iidaka, T., Omori, O., Murata, T., Kosaka, H., Yonekura, Y., Okada, T., Sadato, N., 2001. Neural interaction of the amygdala with the prefrontal and temporal cortices in the processing of facial expressions as revealed by fMRI. J. Cogn. Neurosci. 13, 1035–1047.

Kelso, J.A., Haken, H., 1997. Dynamical Patterns. MIT Press, Cambridge, MA.

Kim, D.-S., Ronen, I., Formisano, E., Ugurbil, K., Mori, S., Goebel, R., 2001. Simultaneous, in vivo visualisation of axonal connectivity and functional activity in cat visual cortex using MRI. Soc. Neurosci. Abstr. 783, 3.

Koch, M.A., Norris, D.G., Hund-Georgiadis, M., 2002. An investigation of functional and anatomical connectivity using magnetic resonance imaging. NeuroImage 16, 241–250.

Maguire, E.A., Mummery, C.J., Büchel, C., 2000. Patterns of hippocampal-cortical interaction dissociate temporal lobe memory subsystems. Hippocampus 10, 475–482.

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Le Goualher, G., Feidler, J., Smith, K., Boomsma, D., Hulshoff Pol, H., Cannon, T., Kawashima, R., Mazoyer, B., 2001. A four-dimensional probabilistic atlas of the human brain. J. Am. Med. Inform. Assoc. 8, 401–430.

McIntosh, A.R., 1999. Mapping cognition to the brain through neural interactions. Memory 7, 523–548.

McIntosh, A.R., Gonzalez-Lima, F., 1994. Structural equation modelling and its application to network analysis in functional brain imaging. Hum. Brain Mapp. 2, 2–22.

Morosan, P., Rademacher, J., Schleicher, A., Amunts, K., Schormann, T., Zilles, K., 2001. Human primary auditory cortex: cytoarchitectonic subdivisions and mapping into a spatial reference system. NeuroImage 13, 684–701.

Paus, T., Jech, R., Thompson, C.J., Comeau, R., Peters, T., Evans, A.C., 1997. Transcranial magnetic stimulation during position emission tomography: a new method for studying connectivity of the human cerebral cortex. J. Neurosci. 17, 3178–3184.

Paus, T., Jech, R., Thompson, C.J., Comeau, R., Peters, T., Evans, A.C., 1998. Dose-dependent reduction of cerebral blood flow during rapid-rate transcranial magnetic stimulation of the human sensorimotor cortex. J. Neurophysiol. 79, 1102–1107.

Rademacher, J., Morosan, P., Schormann, T., Schleicher, A., Werner, C., Freund, H.J., Zilles, K., 2001. Probabilistic mapping and volume measurement of human primary auditory cortex. NeuroImage 13, 669–683.

Repucci, M.A., Schiff, N.D., Victor, J.D., 2001. General strategy for hierarchical decomposition of multivariate time series: implications for temporal lobe seizures. Ann. Biomed. Eng. 29, 1135–1149.

Roland, P., Svensson, G., Lindeberg, T., Risch, T., Baumann, P., Dehmel, A., Frederiksson, J., Halldorson, H., Forsberg, L., Young, J., Zilles, K., 2001. A database generator for human brain imaging. Trends Neurosci. 24, 562–564.

Schormann, T., Zilles, K., 1998. Three-dimensional linear and nonlinear transformations: an integration of light microscopical and MRI data. Hum. Brain Mapp. 6, 339–347.

Schubert, D., Staiger, J.F., Cho, N., Kotter, R., Zilles, K., Luhmann, H.J., 2001. Layer-specific intracolumnar and transcolumnar functional connectivity of layer V pyramidal cells in rat barrel cortex. J. Neurosci. 21, 3580–3592.

Sporns, O., Tononi, G., Edelman, G.M., 2000. Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. Neural Networks 13, 909–922.

Stone, J.V., 2002. Independent component analysis: an introduction. Trends Cogn. Sci. 6, 59–64.

Treves, A., 1993. Mean-field analysis of neuronal spike dynamics. Network 4, 259–284.

Vanduffel, W., Fize, D., Mandeville, J., Nelissen, K., Van Hecke, P., Rosen, B., Tootell, R., Orban, G., 2001. Visual motional processing investigated using contrast agent-enhanced fMRI in awake behaving monkeys. Neuron 32, 565–577.

Zeki, S.M., 1978. Functional specialisation in the visual cortex of the rhesus monkey. Nature 274, 423–428.