Published in final edited form as:

Curr Opin Neurobiol. 2008 December ; 18(6): 582–588. doi:10.1016/j.conb.2008.11.005.

Inferring functional connections between neurons

Ian H. Stevenson¹, James M. Rebesco¹, Lee E. Miller^{1,2}, and Konrad P. Körding^{1,3}

¹Department of Physiology, Feinberg School of Medicine, Northwestern University 303 E. Chicago Ave, Chicago, IL 60611, USA

²Department of Biomedical Engineering, Northwestern University 2145 Sheridan Rd, Evanston, IL 60208, USA

³Department of Physical Medicine and Rehabilitation, Rehabilitation Institute of Chicago 345 East Superior St, Chicago, IL 60611, USA

Summary

A central question in neuroscience is how interactions between neurons give rise to behavior. In many electrophysiological experiments, the activity of a set of neurons is recorded while sensory stimuli or movement tasks are varied. Tools that aim to reveal underlying interactions between neurons from such data can be extremely useful. Traditionally, neuroscientists have studied these interactions using purely descriptive statistics (cross-correlograms or joint peri-stimulus time histograms). However, the interpretation of such data is often difficult, particularly as the number of recorded neurons grows. Recent research suggests that model-based, maximum likelihood methods can improve these analyses. In addition to estimating neural interactions, application of these

© 2008 Elsevier Ltd. All rights reserved.

Correspondence should be addressed to KPK (E-mail: kk@northwestern.edu).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Annotations

•• Truccolo W et al. (2005)

Using data recorded from a population of neurons in primary motor cortex, the authors show that a GLM can improve decoding of hand kinematics. They also present several related models and use an important tool for model selection, not discussed in this review (KS-tests on the time-rescaled spike times).

•• Pillow JW et al. (2008)

Using data from a densely recorded population of retinal ganglion cells, the authors show that a GLM can improve decoding of visual stimuli. The inferred functional connections correspond closely with the known anatomy of the retina, where neighboring cells are highly correlated. Moreover, the authors found that trial-to-trial variation can largely be explained by these functional connections.

• Kass RE et al (2005)

The authors review some of the advanced methods for analyzing spike train data. They thoroughly explain model-based methods and tools to test their validity.

• Pillow J (2007)

The author reviews a number of likelihood-based models of neural coding. Functional connectivity is not the central theme, but maximum likelihood methods and tools for testing their validity (KS-tests on time-rescaled spike times and cross-validation) are thoroughly explained.

• Schneidman E et al (2006)

The authors use a popular alternative to the GLM, a maximum entropy model, to analyze data from a population of retinal ganglion cells. They show that most of the variance in spiking can be explained by pair-wise interactions alone. Higher-order correlations between triplets of neurons, for instance, play a smaller role. They also discuss the “curse of dimensionality,” one of the difficulties in extending descriptive methods to large numbers of neurons.

• Stephan KE et al (2004)

The authors review dynamic causal modeling (DCM), one of the most useful model-based methods for estimating connectivity from fMRI data. Although this is a dynamical model, the power of explaining away and the potentially confounding presence of common input exist here as well.

techniques has improved decoding of external variables, created novel interpretations of existing electrophysiological data, and may provide new insight into how the brain represents information.

The problem of inferring functional connectivity

At a high level, one of the objectives of neuroscience research is to understand how the components of large, complex networks interact. We want to understand how different areas of the brain interact, how groups of neurons within these areas interact, how individual neurons in local circuits interact, and, of course, how all of these interactions relate to the external world. Although these interactions occur on different spatial and temporal scales, in many ways, they reflect a common underlying question: Given recordings of some elements in a network (e.g. groups of neurons, individual neurons, or protein expression), how do we infer and interpret interactions among elements? In the past few decades a number of theoretical and experimental advances have allowed neuroscientists to begin to answer this question for a wide variety of signals ranging from fMRI and PET imaging to simultaneous recordings of many single neurons [1-5]. In this review we focus on the ideas underlying new techniques for the inference of functional connectivity from spike data.

It has long been known that neurons, even when they are far apart in the brain, often do not fire independently from each other but exhibit correlated firing patterns [6]. Dependencies between the firing patterns of observed neurons may be due to a number of different reasons. (1) The neurons may interact monosynaptically or polysynaptically. (2) They may encode similar properties of the environment or planned movements. (3) They may receive common drive from other, unobserved neurons in the nervous system. Algorithms that infer functional connectivity analyze the dependencies among firing patterns and strive to infer how these factors give rise to those dependencies.

For several decades, neurophysiologists have used correlations between neurons to characterize their interactions [3,4,7,8]. Early methods focused on analyzing pairs of neurons using cross-correlograms [6] or joint peri-stimulus time histograms [9]. These methods have become staples of neural data analysis, and have revealed a great deal about the interactions between cortical and subcortical structures [10] and the local interactions in visual [11,12] and auditory cortices [13,14]. These techniques can sometimes reveal the signature of a synaptic connection but provide no well defined way of distinguishing between the kinds of interactions discussed above. Recent developments in model-based, maximum likelihood approaches promise to improve the estimation of connectivity [15,16]. Here we review the statistical principles behind these approaches, some recent applications, and several open questions in neuroscience that these approaches may be able to address.

Explaining Away

Before describing specific models it is helpful to understand how model-based methods address some of the statistical problems that exist with estimating functional connectivity. Consider a situation with three observed neurons (A, B and C), where A excites B and B excites C (Figure 1a). In this situation there will be strong correlations between the activities of all pairs of neurons. However, it is not immediately obvious from the correlations that the interaction between A and C is mediated completely by B -- the cross-correlations seem to suggest a weak interaction between A and C (Figure 1b). Intuitively, if we could somehow control for the influence of B we would find that there is no interaction between A and C (see [17,18], for instance). That is, the interaction between A and C should be explained away when we consider the influence of *both* A and B on C. Model-based methods do exactly this by explicitly modeling all possible connections and fitting the parameters all at once. Figure 1c shows the estimated interactions for this example.

Now consider a second case where neuron A projects to both neurons B and C, and, again, all neurons are observed (Figure 1d). Here, B and C both receive input from A which tends to confound our estimate of the interaction between B and C. Using cross-correlations alone it is difficult for us to determine whether B and C are interacting [19]. Again, model-based methods, by considering all possible interactions at once, allow the correlation between neuron B and C to be explained away by their shared interactions with neuron A (Figure 1f).

While these examples are idealized, they show that correlations between pairs of neurons, in isolation, generally provide an incomplete description of the interactions among observed neurons. In effect, cross-correlation methods describe the probability that a neuron spikes given the recent activity of another neuron. What we often want to describe is the probability that a neuron fires given the recent history of all observed variables, including (but not limited to) the activities of other neurons. Explaining away is a powerful feature of model-based methods. By explicitly modeling all observed variables, many of the potential interactions are discovered to be side-effects of more direct interactions during model fitting.

Unobserved Common Input

When variables in the network are not observed the situation becomes more complicated. For instance, if neuron A in figure 1d is not observed, we probably would infer an interaction between neurons B and C. Such unobserved common-input generally confounds connectivity estimates. In the motor cortex, for instance, movement induces strong input correlations between many neurons. Fortunately, by including movement-related variables in the model, our estimates of functional connectivity may be improved.

Given this caveat, what does it mean for two neurons to be “functionally connected”? One helpful way of thinking about this problem was introduced by Aertsen et al. [20]. They note that it is impossible to uniquely determine the “true” connectivity of a network without recording from all elements. Unobserved elements in the network can always confound connectivity estimates [21]. What we infer is, at best, an approximate description of the network. Given a model and the neurons we observe, the functional connectivity is a reconstruction of the pair-wise connections that best reproduces the data. How well the inferred functional connectivity matches the anatomical connectivity is an issue of future research. Some emerging techniques combine anatomical knowledge with neural signals to further improve inferred functional connectivity [22]. However, even when functional connectivity does not match actual, anatomical connectivity, the results can be useful. The inferred functional connectivity summarizes the statistical relationships governing the interactions between neurons.

Model-based approaches

One of the central ideas in model-based approaches is the idea of a generative model [23]. It is assumed that the observed signals are caused or generated by one or more (potentially hidden) processes. In the case of neurons, it is typically assumed that the firing of each neuron is influenced by the recent activity of some external variables (stimuli or movement), the neuron’s own recent activity, and the recent activity of other observed neurons [24,25]. These influences are assumed to be statistical, for example, when a neuron fires it may increase or decrease the probability of another neuron’s firing. These influences are characterized by parameters such as connection strengths and tuning curve properties. To infer functional connectivity with a model-based approach the parameters of the model are fit so that the probability of the measured neural signals is as high as possible given the parameter values (maximum likelihood) [26]. Within the framework of model-based analysis, different assumptions can be incorporated about the way that observable signals are generated. For instance, spike trains are often modeled

as generalized linear models (GLMs) [27,28]. We will go through these underlying assumptions of this particular model in detail.

The GLM assumes the generative model outlined in Figure 2. It is assumed that the neuron's propensity to fire, called conditional intensity [29], is affected by three factors: the neuron's own recent spiking, the recent spiking of other recorded neurons, and correlated inputs from unobserved populations of neurons that encode external covariates (e.g. stimuli or movement) in some fashion. Using these factors, there are parameters to capture refractory and other intrinsic effects, functional connections with other neurons, and receptive fields or tuning curves (such as the cosine tuning curve in figure 2). The conditional intensity function is calculated as the sum of the three, linearly filtered, factors, passed through a static nonlinearity. Each spike is then randomly drawn according to a Poisson distribution with this rate, and the final output ends up being distributed according to a Cox process [29]. A number of recent studies have used this GLM idea [16,30,31], and there are a number of algorithms that rapidly optimize the parameters of such GLMs [16,32,33].

These methods have been used to analyze results from a range of recent experiments. They have been used to ask how retinal ganglion cells interact [31], and to analyze cultures of neurons *in vitro* [34]. These two cases are of particular importance, since the anatomy of the retina is well known, and the connections in *in vitro* preparations can potentially be imaged. In both the retina and cultures of neurons there is a strong relationship between the spatial layout of the network and the measured functional connectivity. Figure 3 shows an example, from Pillow et al. [31], of the type of results the GLM produces. Model fitting provides estimates of stimulus effects, refractory effects, and the interactions between neurons. Figure 3 shows estimates for a typical ON retinal ganglion cell (A) and a typical OFF retinal ganglion cell (B). The right sections of figure 3A and 3B show the spatial and temporal organization of the input onto the modeled cell. The modeled cell (shaded black in the mosaics) receives positive input from neighboring cells of the same type (gain > 1) and negative input from neighboring cells of the opposite type (gain < 1). The two mosaics (ON and OFF) overlap, but are shown separately here for clarity. Although interactions with non-neighboring neurons (unshaded in the mosaics) are considered during model fitting, they are explained away by the more direct connections between neighbors. This figure demonstrates that the known anatomical properties of the retina are readily reproduced by GLM-based functional connectivity algorithms.

In other cases it is more difficult to observe anatomical connections between individual neurons. GLMs have been used to describe how populations of hippocampal place cells interact [16], and to analyze the properties of neurons in motor cortex [30]. Such analysis is particularly interesting as it promises to improve brain machine interfaces. In brain machine interface applications the objective is to decode the neural signals and estimate the external covariates. When decoding from neural activities, neurons are typically treated as independent given the external covariates. However, generally neurons are *not* independent. Functional connections introduce correlations in the noise structure of these populations and can, thus, degrade the performance of decoding techniques. Functional connectivity methods explicitly model the influence of neurons on one another; by allowing the spurious interactions between neurons to be explained away these methods promise to improve decoding. Indeed, recent research has shown that modeling the interactions between neurons can often significantly improve decoding of an external variable [30,31,35]. Figure 3C, for example, shows the performance of the GLM, as well as several other models, in decoding visual stimuli presented to a population of retinal ganglion cells. Functional connectivity techniques thus are useful beyond simply estimating the way neurons influence each other.

Beyond the basic GLM: Basis Functions and Priors

Several variations and extensions of the basic GLM have been proposed [34-39], and a number of alternative methods have been developed for modeling multi-neuron spike-train data and inferring functional connectivity [15,40-48]. One of the biggest issues that several lines of research have attempted to address, is that maximum likelihood methods, by themselves, often over-fit the data [31,34,37]. By virtue of the fact that these models may contain a large number of parameters, they often describe the original spikes very accurately but fail to generalize well to new sets of spikes. This is generally a problem when many parameters are inferred from a limited set of data. For example, if we record from 100 neurons there are 10,000 possible connections between neurons and, therefore, at least 10,000 parameters that need to be estimated. The number of free parameters can easily be larger than the number of recorded spikes per neuron - making overfitting a central issue.

There are two common ways to reduce such overfitting. Some approaches reduce the number of parameters by assuming that interactions can be described by a specific function or a small number of basis functions [30,31]. For example, instead of fitting 360 parameters for each degree of a directional tuning curve, we can assume that the neuron is cosine tuned and fit one or two parameters. Alternatively, we can incorporate prior knowledge about the nature of the inference problem using Bayes' rule and calculate maximum a posteriori estimates rather than maximum likelihood estimates [26]. For instance, the number of free parameters can be reduced using the assumption of sparse connectivity [31,34,36,37,49] or by assuming that interaction kernels are smooth [37,50]. Prior beliefs have the effect of nonparametrically reducing the effective number of model parameters. These two methods, parameterization and priors, allow estimating functional connectivity from fewer spikes and promise to make the inference of functional connectivity possible for large numbers of recorded neurons.

Discussion

Inference of functional connectivity is a central technique for the analysis of data for imaging signals [51-54] as well as EEG/MEG data [55]. Advanced techniques have been developed that take into account the statistical properties of the signals (such as hemodynamics) [56,57] as well as structural constraints between different brain areas [58,59]. In many ways, these techniques are similar to the techniques we discussed above. Explaining away plays an important role, and unobserved common input is a potential confound. Given the wide use of these techniques for analyzing imaging and electrogram data, it is surprising that similar techniques are not yet in wide use for the analysis of spike data. One reason for the late adoption may be that the statistical properties of spikes (point processes) are somewhat more complicated to formulate statistically. However, we hope that future progress in functional connectivity analysis for spike data may interface productively with results in the imaging and electrogram literature. Analysis of spike data may benefit from the detailed anatomical results from common imaging methods, and analysis of imaging data may be validated by results from spike data which have much higher temporal resolution.

Many basic ideas about connectivity remain to be tested and for the moment some caution should be used in interpreting results from these methods. In cases such as in the hippocampus and in motor cortex it is not yet clear to what extent functional connectivity inferred from spikes mirrors anatomical connectivity or how well functional connectivity generalizes from one task to the next. To what extent is functional connectivity a real property of the nervous system instead of a statistical property of the algorithm and the task? Although there are still many questions that need to be asked about the methods, these approaches certainly promise to be useful as they allow us to ask how neurons interact, a central question in neuroscience. They will allow us to ask a whole range of further questions in the future. How much do higher-

order interactions and common-input matter? How do functional connections evolve over time, during adaptation, learning, sleep, or aging? How do local functional connections differ across brain regions? And can we use electrical stimulation or drugs to change connectivity?

Lastly, it is important to mention that functional connectivity techniques should be combined with other approaches to brain science. Anatomical connections set constraints that rule out certain functional connections. Synaptic physiology rules out certain interaction kernels. Ultimately algorithms for the inference of functional connectivity should draw on advances in all these areas to improve the resulting estimates.

References

1. Friedman N. Inferring Cellular Networks Using Probabilistic Graphical Models. *Science* 2004;303:799–805. [PubMed: 14764868]
2. Yoshimura Y, Dantzker JL, Callaway EM. Excitatory cortical neurons form fine-scale functional networks. *Nature* 2005;433:868–873. [PubMed: 15729343]
3. Brown EN, Kass RE, Mitra PP. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience* 2004;7:456–461.
4. Kass RE, Ventura V, Brown EN. Statistical Issues in the Analysis of Neuronal Data. *Journal of Neurophysiology* 2005;94:8–25. [PubMed: 15985692]
5. Sporns O, Chialvo DR, Kaiser M, Hilgetag CC. Organization, development and function of complex brain networks. *Trends in Cognitive Sciences* 2004;8:418–425. [PubMed: 15350243]
6. Perkel DH, Gerstein GL, Moore GP. Neuronal Spike Trains and Stochastic Point Processes: II. Simultaneous Spike Trains. *Biophysical Journal* 1967;7:419. [PubMed: 4292792]
7. Dayan, P.; Abbott, LF. Theoretical neuroscience: computational and mathematical modeling of neural systems. MIT Press; 2001.
8. Rieke, F.; Warland, D.; van Steveninck, RR.; Bialek, W. Spikes: exploring the neural code. MIT Press; 1999.
9. Gerstein GL, Perkel DH. Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation. *Science* 1969;164:828–830. [PubMed: 5767782]
10. Sillito AM, Jones HE, Gerstein GL, West DC. Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature* 1994;369:479–482. [PubMed: 8202137]
11. Ts'o DY, Gilbert CD, Wiesel TN. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience* 1986;6:1160. [PubMed: 3701413]
12. Gochin PM, Miller EK, Gross CG, Gerstein GL. Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Experimental Brain Research* 1991;84:505–516.
13. Dickson JW, Gerstein GL. Interactions between neurons in auditory cortex of the cat. *Journal of Neurophysiology* 1974;37:1239–1261. [PubMed: 4436698]
14. Espinosa IE, Gerstein GL. Cortical auditory neuron interactions during presentation of 3-tone sequences: effective connectivity. *Brain Res* 1988;450:39–50. [PubMed: 3401720]
15. Paninski L, Pillow JW, Simoncelli EP. Maximum Likelihood Estimation of a Stochastic Integrate-and-Fire Neural Encoding Model. *Neural Computation* 2004;16:2533–2561. [PubMed: 15516273]
16. Okatan M, Wilson MA, Brown EN. Analyzing Functional Connectivity Using a Network Likelihood Model of Ensemble Neural Spiking Activity. *Neural Computation* 2005;17:1927–1961. [PubMed: 15992486]
17. Brody CD. Correlations Without Synchrony. *Neural Computation* 1999;11:1537–1551. [PubMed: 10490937]
18. Brody CD. Disambiguating Different Covariation Types. *Neural Computation* 1999;11:1527–1535. [PubMed: 10490936]
19. Moore GP, Segundo JP, Perkel DH, Levitan H. Statistical Signs of Synaptic Interaction in Neurons. *Biophysical Journal* 1970;10:876. [PubMed: 4322240]

20. Aertsen AM, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of "effective connectivity". *Journal of Neurophysiology* 1989;61:900–917. [PubMed: 2723733]
21. Horwitz B. The elusive concept of brain connectivity. *Neuroimage* 2003;19:466–470. [PubMed: 12814595]
22. McIntosh AR, Gonzalez-Lima F. Structural equation modeling and its application to network analysis in functional brain imaging. *Human Brain Mapping* 1994;2:2–22.
23. Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press; 2000.
24. Pillow, J. Likelihood-Based Approaches to Modeling the Neural Code. In: Kenji Doya, SI.; Pouget, Alexandre; Rao, Rajesh P.N., editors. *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press; 2007. p. 53-70.
25. Simoncelli, EP.; Paninski, L.; Pillow, J.; Schwartz, O. Characterization of neural responses with stochastic stimuli. In: Gazzaniga, M., editor. *The New Cognitive Neurosciences*. Vol. edn 3. MIT Press; 2004.
26. MacKay, DJC. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press; 2003.
27. McCullagh, P.; Nelder, JA. *Generalized Linear Models*. Vol. 2nd ed.. Chapman and Hill; London: 1989.
28. McCulloch CE. *Generalized Linear Models*. *Journal of the American Statistical Association* 2000;95
29. Daley, DJ.; Vere-Jones, D. *An Introduction to the Theory of Point Processes*. Springer; 2003.
30. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology* 2005;93:1074–1089. [PubMed: 15356183]
31. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008
32. Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics* 1988;59:189–200. [PubMed: 3179344]
33. Chornoboy ES, Schramm LP, Karr AF. Maximum likelihood identification of neural point process systems. *Biological Cybernetics* 1988;59:265–275. [PubMed: 3196770]
34. Rigat F, de Gunst M, van Pelt J. Bayesian modelling and analysis of spatio-temporal neuronal networks. *Bayesian Analysis* 2006;1:733–764.
35. Kulkarni JE. Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems* 2007;18:375–407.
36. Gerwinn S, Macke JH, Seeger M, Bethge M. Bayesian Inference for Spiking Neuron Models with a Sparsity Prior. *Advances in Neural Information Processing Systems* 2007;20
37. Stevenson I, Rebesco J, Hatsopoulos N, Haga Z, Miller L, Körding K. Bayesian inference of functional connectivity and network structure from spikes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* to appear.
38. Eldawlatly S, Jin R, Oweiss K. Identifying functional connectivity in large scale neural ensemble recordings: A multiscale data mining approach. *Neural Computation* to appear.
39. Czanner G, Eden UT, Wirth S, Yanike M, Suzuki WA, Brown EN. Analysis of between-trial and within-trial neural spiking dynamics. *J Neurophysiol* 2008;99:2672–2693. [PubMed: 18216233]
40. Utikal KJ. A new method for detecting neural interconnectivity. *Biological Cybernetics* 1997;76:459–470.
41. Smith AC, Brown EN. Estimating a State-Space Model from Point Process Observations. *Neural Computation* 2003;15:965–991. [PubMed: 12803953]
42. Nykamp DQ. Revealing Pairwise Coupling in Linear-Nonlinear Networks. *SIAM Journal on Applied Mathematics* 2005;65:2005–2032.
43. Martignon L, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E. Neural Coding: Higher-Order Temporal Patterns in the Neurostatistics of Cell Assemblies. *Neural Computation* 2000;12:2621–2653. [PubMed: 11110130]
44. Schneidman E, Berry MJ, Ii RS, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 2006;440:1007. [PubMed: 16625187]

45. Seth AK, Edelman GM. Distinguishing Causal Interactions in Neural Populations. *Neural Computation* 2007;19:910–933. [PubMed: 17348767]
46. Baccalá LA, Sameshima K. Partial directed coherence: a new concept in neural structure determination. *Biological Cybernetics* 2001;84:463–474. [PubMed: 11417058]
47. Matthias W, Björn S, Wolfram H, Karin S, Lutz L, Daniel K, Reinhard B, Jens T, Herbert W. Comparison of linear signal processing techniques to infer directed interactions in multivariate neural systems. *Signal Process* 2005;85:2137–2160.
48. Schreiber T. Measuring Information Transfer. *Physical Review Letters* 2000;85:461. [PubMed: 10991308]
49. Paninski L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Comput. Neural Syst* 2004;15:243–262.
50. Sahani, M.; Linden, JF. Evidence Optimization Techniques for Estimating Stimulus-Response Functions; *Advances in Neural Information Processing Systems 15: Proceedings of the 2002 Conference*; 2003;
51. Stephan KE, Harrison LM, Kiebel SJ, David O, Penny WD, Friston KJ. Dynamic causal models of neural system dynamics: current state and future extensions. *Journal of Biosciences* 2007;32:129–144. [PubMed: 17426386]
52. Jirsa, VK.; McIntosh, AR. Handbook of Brain Connectivity. In: Jirsa, Viktor K.; McIntosh, AR., editors. *Handbook of Brain Connectivity*. Springer; Berlin: 2007. 2007
53. Weiler N, Wood L, Yu J, Solla SA, Shepherd GMG. Top-down laminar organization of the excitatory network in motor cortex. *Nat Neurosci* 2008;11:360–366. [PubMed: 18246064]
54. Hampson M, Peterson BS, Skudlarski P, Gatenby JC, Gore JC. Detection of functional connectivity using temporal correlations in MR images. *Human Brain Mapping* 2002;15:247–262. [PubMed: 11835612]
55. Takahashi DY, Baccal LA, Sameshima K. Connectivity Inference between Neural Structures via Partial Directed Coherence. *Journal of Applied Statistics* 2007;34:1259–1273.
56. Stephan KE, Harrison LM, Penny WD, Friston KJ. Biophysical models of fMRI responses. *Current Opinion in Neurobiology* 2004;14:629–635. [PubMed: 15464897]
57. Harrison LM, David O, Friston KJ, Philosophical transactions-Royal Society of London. Stochastic models of neuronal dynamics. *Biological sciences* 2005;360:1075–1091. [PubMed: 16087449]
58. Friston K, Harrison L, Daunizeau J, Kiebel S, Phillips C, Trujillo-Barreto N, Henson R, Flandin G, Mattout J. Multiple sparse priors for the M/EEG inverse problem. *Neuroimage*. 2007
59. Harrison LM, Penny W, Daunizeau J, Friston KJ. Diffusion-based spatial priors for functional magnetic resonance images. *Neuroimage*. 2008

Figure 1.

Two simulations of explaining away. In the first case (a), although cross-correlations indicate that there may be an interaction between A and C (gold box in panel b), it is explained away in model-based estimates of functional connectivity (c). In the second case (d), cross-correlations suggest that there may be a weak interaction between B and C (gold box in panel e), but this interaction is explained away by the fact that both neurons receive input from A (f).

Figure 2.

The Generalized Linear Model. For each neuron the instantaneous firing rate may depend on external covariates (the direction of hand movement, for instance), the neuron's recent spike history, and the recent spike history of other recorded neurons. These variables are linearly filtered by tuning curves, refractory effects, and functional connections, and the firing rate is modeled as the sum of these effects passed through a static nonlinearity (e.g. an exponential). We can fit the parameters of the model (tuning curves, refractory effects, and functional connections) by comparing the predicted spiking to the actual spikes.

Figure 3.

Results from Pillow et al. (2008). The GLM fitted to an example ON retinal ganglion cell (A) and an example OFF retinal ganglion cell (B). Each fit contains a stimulus filter (both temporal and spatial components), refractory effects, and interactions (functional connections) with other neurons. The ON and OFF mosaics (panels A and B, upper right) show the spatial organization of the interactions; the modeled cell is shaded in black. Note that the functional connections have a strong spatial organization - cells receive positive input from neighboring cells of the same type (red lines in A, blue lines in B) and negative input from neighboring cells of the opposite type (blue lines in A, red lines in B). The two mosaics (ON and OFF) overlap, but are shown separately for clarity. (C) shows the decoding accuracy for the population using a linear model, a Poisson model (no refractory effects or functional connections), an uncoupled GLM (no functional connections), and the full GLM. Adapted by permission from Macmillan Publishers Ltd: Nature, copyright (2008).