Bionet Group

In Silico and In Vivo Information Representation and Neural Computation
http://bionet.ee.columbia.edu

CNS*2009 Workshop on

Methods of Information Theory in Computational Neuroscience

Wednesday and Thursday, July 22-23, 2009

Overview

Methods originally developed in Information Theory have found wide applicability in computational neuroscience. Beyond these original methods there is a need to develop novel tools and approaches that are driven by problems arising in neuroscience.

A number of researchers in computational/systems neuroscience and in information/communication theory are investigating problems of information representation and processing. While the goals are often the same, these researchers bring different perspectives and points of view to a common set of neuroscience problems. Often they participate in different fora and their interaction is limited.

The goal of the workshop is to bring some of these researchers together to discuss challenges posed by neuroscience and to exchange ideas and present their latest work.

The workshop is targeted towards computational and systems neuroscientists with interest in methods of information theory as well as information/communication theorists with interest in neuroscience.

References

Baddeley, Hancock & Foldiak (eds), Information Theory and the Brain, Cambridge University Press, Cambridge, UK, 2000.

Borst & Theunissen, Information Theory and Neural Coding, Nature Neuroscience, 2, pp. 947 - 957, 1999. doi:10.1038/14731

D.H. Johnson, Dialogue Concerning Neural Coding and Information Theory, August 2003.

C.E. Shannon, A Mathematical Theory of Communication, Bell System Technical Journal, vol. 27, pp. 379-423 and 623-656, 1948.

Journal of Computational Neuroscience

Call for Papers

Special issue on Methods of Information Theory in Neuroscience Research

Alexander G. Dimitrov , Aurel A. Lazar and Jonathan D. Victor, (editors). Submission deadline: December 15, 2009.

Methods originally developed in Information Theory have found wide applicability in computational neuroscience. Beyond these original methods, novel tools and approaches have been developed that are driven by problems arising in neuroscience. A number of researchers in computational/systems neuroscience and in information/communication theory are investigating problems of information representation and processing.

The goal of the special issue of the Journal of Computational Neuroscience is to showcase the latest techniques, approaches and results in this area. The subject of the papers must fit the Aims & Scope of the journal, and must in particular not be purely methodological but also illustrate results that advance our understanding of brain function in a broad sense. The papers must contain new material, but we encourage the authors to also include review material to help the reader fully understand the context of the study. Papers that include experimental work are especially encouraged.

Submission is open to all as long as it fits the above criteria. The papers will go through the standard review process, with the same criteria as for normal articles submitted to the journal. They will be further reviewed for relevance to the special issue by the guest editors listed above. At submission, please indicate in comments that you wish to be considered for this issue.

Organizers

Aurel A. Lazar, Department of Electrical Engineering, Columbia University and Alex G. Dimitrov, Center for Computational Biology, Montana State University.

Invited Speakers

Bayesian Population Decoding of Spiking Neurons

Matthias Bethge, Max Planck Institute for Biological Cybernetics, Tuebingen.

The timing of action potentials in spiking neurons depends on the temporal dynamics of their inputs, and contains information about temporal fluctuations of a stimulus. Leaky integrate-and-fire neurons constitute a popular class of encoding models in which spike times depend directly on the temporal structure of their inputs. However, optimal decoding rules for these models have only been studied explicitly in the noiseless case. Here, we study decoding rules for probabilistic inference of a continuous stimulus from the spike times of a population of leaky integrate-and-fire neurons with noise threshold fluctuations. We derive three algorithms for approximating the posterior distribution over stimuli as a function of the observed spike trains. In addition to a reconstruction of the stimulus, we thus obtain an estimate of the uncertainty as well. Furthermore, we also derive a `spike-by-spike' online decoding scheme that recursively updates the posterior with the arrival of each new spike. We use these decoding rules to reconstruct time varying stimuli represented by a Gaussian process from spike trains of single neurons as well as neural populations.

Joint work with Sebastian Gerwinn and Jakob Macke.

Informational Robustness of the Hippocampal Spatial Map: Topological Analysis

Yuri Dabaghian (1), Facundo Memoli (2), Gurjeet Singh (2), Gunnar Carlsson (2), Loren Frank (1),
(1) Department of Physiology, University of California at San Francisco, CA. (2) Department of Mathematics, Stanford University, CA.

It is well known that the hippocampus plays crucial role in creating a spatial representation of the environment and in forming spatial memories. Each active hippocampal neuron (a place cell) tends to fire in its “place field’’ – a restricted region of the animal’s environment, so that the ensemble of active place cells encodes a spatial map. However the exact informational contents of this map are currently not clear. The hippocampal map may represent the connectivity of locations in the environment, i.e. be a topological map, or it may contain information about distances and angles and hence be more geometric in nature. The latter possibility is supported by the majority of current theories, which suggest that the hippocampus explicitly represents geometric elements of space derived from a path integration process that takes into account distances and angles of self motion information.

Several recent experiments indicate that the sequential structure of the hippocampal spatial map remains invariant with respect to a significant range of geometrical transformations of the environment [1,2], which indicates that the temporal ordering of spiking from hippocampal neural ensembles is the key determinant of the spatial information communicated to downstream neurons. This implies that the map itself is better understood as representing the topology of the animal's environment and suggests that the actual role of the hippocampus is to encode topological memory maps, where the patterns of ongoing neural activity represent the connectivity of locations in the environment or the connectivity of elements of a memory trace. This hypothesis can be tested and studied both theoretically and experimentally. Specifically, the topological approach allows creating computational models that impose particular relationships on the parameters of the hippocampal map and predict some of its properties, such as the parameter ranges for informational stability of the map, which must agree with the experimentally observed parameters of firing activity in the hippocampus.

Hence we investigate the robustness of simulated hippocampal topological maps of different geometric complexity and dimensionality with respect to independent variations of place cell activity parameters, such as the distributions of the firing rate, the of sizes of place fields and the number of cells in the map, using the Persistent Homology method [2]. We find the maximal range of topological stability for each parameter independently and study the relationships between the parameters that ensure the stability of the map and then compare the results with the experimentally observed values and the dynamics of firing activity and find that our theoretical framework is consistent with experimental data. This provides a possibility to understand some of the physiological features of the hippocampal map based on the informational analysis of its firing activity.

Acknowledgments

This work was supported by NIH grants F32-NS054425-01 and MH080283, by the Sloan and Swartz Foundations, and by DARPA grants HR0011-05-1-0007 and FA8650-06-1-7629.

References

[1] K. Gothard, W. Skaggs, K. Moore, B. McNaughton, J. of Neuroscience, 16(2) (1996).
[2] K. Diba and G. Buzsáki, J. of Neuroscience, 28, 13448-13456 (2008).
[3] H. Edelsbrunner, D. Letscher and A. Zomorodian, Simplification, FOCS 2000: 454-463.

Soft Clustering Decoding of Neural Codes

Alexander G. Dimitrov, Center for Computational Biology, Montana State University.

Methods based on Rate Distortion theory have been successfully used to cluster stimuli and neural responses in order to study neural codes at a level of detail supported by the amount of available data. They approximate the joint stimulus-response distribution by soft-clustering paired stimulus-response observations into smaller reproductions of the stimulus and response spaces. An optimal soft clustering is found by maximizing an information-theoretic cost function subject to both equality and inequality constraints, in hundreds to thousands of dimensions.

The method of annealing has been used to solve the corresponding high dimensional non-linear optimization problems. The annealing solutions undergo a series of bifurcations in order to reach the optimum, which we study using bifurcation theory in the presence of symmetries. The optimal models found by the Information Distortion methods have symmetries: any classification of the data can lead to another equivalent model simply by permuting the labels of the classes. These symmetries are described by SN, the algebraic group of all permutations on N symbols. The symmetry of the bifurcating solutions is dictated by the subgroup structure of SN. In this contribution we describe these symmetry breaking bifurcations in detail, and indicate some of the consequences of the form bifurcations.

Synaptic Encoding

Aurel A. Lazar, Department of Electrical Engineering, Columbia University.

Information Theoretic Analysis of Population Coding of Auditory Space

Nicholas A. Lesica, Department of Biology, Ludwig-Maximilians University, Munich.

With advances in electrophysiology and imaging, it is becoming increasingly common to study the joint responses of neuronal populations. Calculating mutual information directly from the experimental responses of large populations may not be possible, but the responses can be decoded and the mutual information between the actual and decoded stimulus can be measured (Rolls et al., J. Neurophysiol, 1998). This approach has many benefits, the most important of which is that the dimensionality of the 'response' space depends only on the number of distinct stimuli. We used a decoding approach to investigate the coding of auditory space in gerbils and barn owls. We investigated the robustness of different measures of decoder performance: the number of decoding errors, the mutual information measured from the decoded responses (which is affected by both the number and the distribution of decoding errors), and the mutual information measured from the 'distances' between responses (which can be thought of as decoder confidence). Our results demonstrate the utility of information theoretic measures for the analysis of large populations with different strategies for coding auditory space.

Information Theoretic Approaches to Predicting Optimal Neuronal Stimulus-Response Curves

Mark McDonnell, Institute for Telecommunications Research, University of South Australia, Adelaide.

Estimating Mutual Information of Multivariate Spike Count Distributions Using Copulas

Arno Onken (1,2), Steffen Grünewälder (1), and Klaus Obermayer (1,2)
(1) Technische Universität Berlin; (2) Bernstein Center for Computational Neuroscience, Berlin

Neural coding is typically analyzed by applying information theoretic measures like the mutual information between stimuli and spike count responses for a given bin size. In order to estimate the mutual information, a multivariate noise model of the spike counts is required. Here, we use copulas to construct discrete multivariate distributions that are appropriate to model dependent spike counts of several neurons. With copulas it is possible to use arbitrary marginal distributions such as Poisson or negative binomial that are better suited for modeling single neuron noise distributions than the most often applied normal approximation. Furthermore, copulas place a wide range of dependence structures at the disposal and can be used to analyze higher order interactions. We present a framework for model fitting of copula based distributions and apply Monte-Carlo techniques to estimate entropy and mutual information.

Sheila Nirenberg, Department of Physiology and Biophysics, Weill Medical College, Cornell University.

Higher-Order Correlations in Large Neuronal Populations

Benjamin Staude, Sonja Grün and Stefan Rotter, Bernstein Center for Computational Neuroscience, Freiburg.

Spiking neurons are known to be quite sensitive for the higher-order correlation structure of their respective input populations (Kuhn et al. 2003). What is the role of these correlations in cortical information processing?

A prerequisite to answering this question is an appropriate framework to describe and effectively estimate the correlation structure of neuronal populations. Approaches available thus far suffer from the combinatorial explosion of the number of parameters that grows exponentially with the number of recorded neurons. As a consequence, methods that go beyond pairwise correlations and aim for estimating genuine higher-order effects require vast samples, rendering them essentially inapplicable to populations of more than ~10 neurons.

Here, we discuss the compound Poisson process as an intuitive and flexible model for correlated populations of spiking neurons. Based on this generative model, we present novel estimation techniques to infer the correlation structure of a neural population from sampled spike trains (Ehm et al. 2007; Staude et al. 2009). Our techniques can provide conclusive evidence for higher-order correlations in rather large populations of ~100 neurons, based on sample sizes that are compatible with current physiological in vivo recording technology.

Kuhn A, Aertsen A, Rotter S. Higher-order statistics of input ensembles and the response of simple model neurons. Neural Computation 15(1): 67-101, 2003.
Ehm W, Staude B, Rotter S. Decomposition of neuronal assembly activity via empirical de-Poissonization. Electronic Journal of Statistics 1: 473-495, 2007.
Staude B, Rotter S, Grün S. CuBIC: cumulant based inference of higher-order correlations in massively parallel spike trains. Under review.

GELO - a Simple Stochastic Oscillator Traces Cellular and Coding Mechanisms

Gaby Schneider, Department of Informatics and Mathematics, University of Frankfurt.

Temporal coordination of neuronal spike trains, such as synchronous activity, oscillation, repetitive bursting or systematic temporal delays, have been hypothesized to be relevant for information processing. We present a simple stochastic model called the GELO which can describe and analyze these features in one theoretical framework.

Regular pacemakers and processes with repetitive bursts are described with the same set of parameters. The regularity of the oscillation and the degree of burstiness are estimated by fitting the model to the autocorrelation histogram. This allows a reliable burst detection in individual spike trains even in nonstationary conditions when burst surprise measures can fail. For parallel processes with a joint oscillation, the model can capture fine temporal structure across the units. By relating the auto correlation to the cross correlation histogram, it can also measure the degree to which all units share the same underlying oscillation. We illustrate the method with application to single spike trains recorded in the substantia nigra of mice and to a set of parallel spike-trains obtained from the visual cortex of the cat.

Acknowledgements: This is joint work with Markus Bingmer (Inst. of Mathematics, Frankfurt University), Felipe Gerhard and Danko Nikolic (FIAS, Frankfurt University), Jochen Roeper and Julia Schiemann (Institute of Neurophysiology, Frankfurt University). This work was supported by the German Bundesministerium fuer Bildung und Forschung (Bernstein Focus Neurotechnology, Frankfurt)

Speech Recognition Using Spiking Neural Networks

Benjamin Schrauwen, Electronics and Information Systems Department, Ghent University.

Spiking Neural Networks are typically used as mere modeling tools in neuroscience. It has however recently been shown that Spiking Neurons are (1) computationally more efficient than analog neurons, (2) can intrinsically process temporal signals, (3) are very suited for both hardware as well as software implementation, (4) and are biologically more realistic. We will demonstrate these various properties on an isolated digits speech recognition application. One of our main research goals is to research all aspects needed to turn Spiking Neural Networks in a tool that can be relatively easily used in an engineer's toolbox. All issues from spike train encoding, over various learning approaches for Spiking Neural Networks, to their efficient hardware and software implementations will be touched upon.

Maximum Entropy Decoding of Multivariate Neural Spike Trains

Simon R. Schultz, Department of Bioengineering, Imperial College.

Distinguishing between Proximal and Distal Stimuli in Sensory Coding

Lars Schwabe, Department of Computer Science and Electrical Engineering, University of Rostock.

Theoretical investigations of sensory coding are often based on optimality assumptions (like efficient coding), possibly with added constraints on desirable properties of the neural code (like sparseness or a factorial code). Most theoretical investigations and experiments, however, neglect a distinction well known in psychophysics, namely the distinction between the proximal and the distal stimuli. In the context of sensory coding, the distal stimuli can be thought of as the “relevant” aspect of the environment (like the 3D properties of objects), whereas the proximal stimuli correspond to the signals sensed at the sensory periphery (like their 2D projection in the visual system). I will briefly recapitulate some theoretical and experimental approaches to sensory coding and consider most of them as encoding the proximal stimuli. Then, I will apply the Information Bottleneck method (Tishby et al., 1999) in order to derive compressed representations of proximal stimuli, such that the compressed representation conveys maximal information about the corresponding distal stimuli. This approach can be considered as an implementation of Purves empirical theory of perception (Purves and Lotto, 2003), and I will illustrate it using examples from the visual and vestibular system. Ideas for the crucial and most challenging aspect of this approach will be presented and shall be discussed: How does the nervous system learn such representations without direct access to the distal stimuli?

Stefano Panzeri, The Italian Institute of Technology, Genoa.

Raul Vicente, FIAS, Frankfurt am Main.

Jonathan D. Victor, Department of Physiology and Biophysics, Weill Medical College of Cornell University.

Cliques in Neural Ensembles and Co-Evolutionary Learning in Liquid Architectures

Yehoshua (Josh) Y. Zeevi, Department of Electrical Engineering, Technion.

Joint work with Igal Raichelgauz and Karina Odinaev.
