II-32. Multimap formation in the Visual Cortex

Rishabh Jain RISHABH@USC.EDU
Bartlett Mel MEL@USC.EDU

University of Southern California

Cells in V4 and IT respond selectively to a wide variety of moderately complex object features [1], [2], perhaps numbering in the hundreds or thousands or more. Given that object recognition is location-invariant, it follows that, (1) every location in the visual field is analyzed in parallel by a large number of different feature types, which implies that (2) every feature type forms its own retinotopic map coarsely tiling the visual field. The physical interdigitation of neurons participating in a potentially large number of distinct feature maps covering the visual field, each of which must develop separately and unsupervised, presents a significant challenge for conventional map-formation algorithms. In particular, within such a “multimap”, a neuron is surrounded by other neurons coding many feature types different from its own and from each other (see [4] [5] and [3] for data suggestive of this), preventing neurons from co-training, i.e. sharing sensory data, within cohorts defined by a purely spatial criterion. Co-training during development must therefore depend both on the spatial proximity of two neurons as well as their functional similarity. In this work, we have developed a new self-organizing map formation algorithm derived from a Kohonen-style SOM but that uses a hybrid spatial-functional similarity criteria for determining which neurons wire and learn together. Preliminary results show that pure spatial algorithms confronted with multi-feature-type data tend to produce a patchwork of relatively large regions of tissue containing neurons of a single feature type (Figure panel A), whereas the hybrid learning rule produces multiple, finely interdigitated smooth feature maps in which neurons are in general surrounded by cells of other types (Figure panel B). Multimap learning algorithms also have the potential to help explain species-specific differences in the fine-scale smoothness vs. roughness of feature maps in primary sensory areas [4].

doi:

II-33. Extracting rotational structure from motor cortical data

John Cunningham[1] JPC74@CAM.AC.UK
Mark Churchland[2] CHURCH@STANFORD.EDU
Matthew Kaufman[2] MATT235@STANFORD.EDU
Krishna V. Shenoy[2] SHENOY@STANFORD.EDU

[1]University of Cambridge
[2]Stanford University

Dimensionality reduction techniques such as PCA are a cornerstone of analyzing high dimensional data. PCA involves an eigenvalue decomposition of the data covariance matrix, which produces ranked orthogonal dimensions that can be used to linearly project the data to lower dimension. The covariance matrix is but one choice of summary matrix that can be used for linear projections. When the data are time series, the same decomposition can be used on a linear description of the dynamics (instead of the data covariance) to obtain a different projection. This ‘dynamical PCA’ produces projections representing the largest eigenvalues of the linear dynamical system. Linear dynamical systems capture scaling and rotational aspects of the data (and combining scalings/rotations yields familiar features such as shear, projection, reflection, etc.). Dynamical PCA makes no distinction between scaling and rotation, seeking only directions of largest and most consistent dynamical activity. However, in some settings, one may hypothesize the existence of certain types of dynamics, and thus an algorithm is desired that can verify those dynamics. For example, it is common for neural circuits to generate oscillations (or rotations), and we seek projections that best capture these fundamental aspects of the neural response. Here we introduce jPCA, which specifically extracts projections of largest *rotational* dynamics. Using skew-symmetric matrices, the jPCA algorithm is an extension of fitting a linear dynamical system. It has a unique, closed-form solution that can be quickly computed. Importantly, like PCA and dynamical PCA, jPCA produces simple projection vectors - orthogonal linear directions - and so the interpretation of jPCA is identical to PCA or any other linear projection of the data. We motivate and present details of the method, and we discuss its importance in extracting rotational structure from electrophysiological data recorded from primate motor cortex.

doi:

II-34. Fast low-SNR high-dimensional optimal filtering, applied to inference of dynamic receptive fields

Liam Paninski LIAM@STAT.COLUMBIA.EDU
Kamiar Rahnama Rad KAMIAR@STAT.COLUMBIA.EDU
Jonathan Huggins JHH2143@COLUMBIA.EDU
Columbia University

Kalman filtering is a workhorse of statistical time series analysis: it is computationally tractable in many real-world settings and implements the optimal Bayesian filter in the linear-Gaussian setting. However, the state variable in many problems is very high-dimensional. Standard implementations of the Kalman filter require $O(N^3)$ time and $O(N^2)$ space, where N is the dimensionality of the state variable, and are therefore impractical. In this paper we note that if a relatively small number of low-SNR observations are available per time step, the Kalman equations may be approximated in terms of a low-rank perturbation of the steady-state (zero-SNR) solution. In many cases this approximation may be computed and updated very efficiently (often in just $O(N)$ or $O(N \log N)$ time and space), using fast methods from numerical linear algebra. This opens up possibility of real-time adaptive experimental design and optimal control in systems of much larger dimensionality than was previously possible. We detail an application to smoothing high-dimensional neuroscience data.

doi:

II-35. Stimulus Evoked Responses in Spontaneously Active Two-Population Neural Networks

Kanaka Rajan KRAJAN@PRINCETON.EDU
Princeton University

To make optimal use of experimental data on the anatomy and physiology of cortical circuits and to account for the effects of plasticity and neuromodulation, we must understand the relationship between the synaptic and neuronal properties of a network, and the activity that it produces. Given detailed knowledge of the properties of a complex network, can we predict what it will do? If we know how plasticity mechanisms or modulators change those properties, can we predict how the activity will change? In this work, we develop a mean-field theory to analyze input-driven responses in networks of neurons with changing connectivity patterns – from uniform synapses to synapses drawn from different probability distributions. We study stimulus-evoked responses in model networks in which only a fraction of the neurons within the network receive the external drive directly. We measure the signal power in the neurons linked to the input only through recurrent polysynaptic connections in the network. This power has a unique dependence on different synaptic parameters and as a result, on the different activity patterns in the network (as a specific example, the synaptic variance at which the signal power is maximally amplified corresponds to chaotic spontaneous activity in the network, but strengthening the mean synaptic strength does not have the same effect.) Generally as networks progress from weak to stronger synapses, we find that stimulus-evoked responses describe a trajectory in the two-dimensional space composed of the mean and the variance of synaptic strengths. The shape and orientation of the trajectory consequently informs us of the regimes in which meaningful responses can be most efficiently extracted from network activity, and how this feature changes as a function of the spatiotemporal properties of the stimulus.

doi: