UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

BERKELEY • DAVIS • IRVINE • LOS ANGELES • RIVERSIDE • SAN DIEGO • SAN FRANCISCO SANTA BARBARA • SANTA CRUZ

KENNETH D. MILLER, PH.D.
PROFESSOR
DEPTS. OF PHYSIOLOGY AND OTOLARYNGOLOGY
NEUROSCIENCE, BIOENGINEERING AND BIOPHYSICS GRADUATE PROGRAMS
W.M. KECK FOUNDATION CENTER FOR INTEGRATIVE NEUROSCIENCE
CO-DIRECTOR, SLOAN CENTER FOR THEORETICAL NEUROBIOLOGY

MAIL: DEPT. OF PHYSIOLOGY
UNIVERSITY OF CALIFORNIA
513 PARNASSUS
SF, CA 94143-0444
TEL: (415) 476 - 8217
FAX: (415) 476 - 4929
EMAIL: ken@phy.ucsf.edu
WWW: http://www.keck.ucsf.edu/~ken

November 5, 2003

I am writing to strongly recommend Tatyana Sharpee for a faculty position in a field, such as biological physics or theoretical neurobiology, consistent with her interdisciplinary training and work. Dr. Sharpee is a brilliant, extremely focused, and extremely hardworking scientist who is bringing her training from theoretical physics to illuminate difficult and fundamental theoretical issues in neuroscience. (Because this letter will be read by committees in a variety of disciplines, I will be somewhat pedantic in describing Tatyana's work and will generally not assume basic neuroscience or mathematical knowledge; I ask the forbearance of the more knowledgeable readers.)

After very successful graduate work in theoretical physics, Tatyana came to UCSF as a Fellow of our Sloan Center for Theoeretical Neurobiology in the Fall of 2001. The Sloan Center is an extremely selective program providing two years of funding for outstanding Ph.D.'s in theoretical fields, such as physics, mathematics, or computer science, to gain training to begin work in the field of neuroscience. Tatyana was an outstanding applicant, one of two admitted her year out of around 60 applicants, and she has been an absolutely outstanding fellow since she arrived. While most of our fellows take a long while to get up to speed in neuroscience before becoming productive, Tatyana has been wonderfully productive from the very beginning. During her first month as a fellow she took the Woods Hole course in computational neuroscience directed by Bill Bialek and Rob de Ruyter, and in that time, working with Bill, she developed a novel solution to a fundamental problem of neuroscience. Since then she has focused on developing and implementing that solution, applying it to real neural data as well as to model data, and learning to do the experiments herself to acquire such data.

The problem Tatyana is addressing is that of determining, from first principles, the stimulus attributes to which a neuron is sensitive. The problem can be characterized as follows. To characterize an area of the brain, one typically inserts electrodes in the area and isolates the action potentials or "spikes" of a single neuron (action potentials are all-or-none voltage events that transmit signals to other neurons, initiated roughly when a threshold voltage is reached; in most parts of the brain, subthreshold voltage modulations do not result in signals to other neurons but instead are "private" to the given neuron, so action potentials are the "language" of most of the brain). One then attempts to determine what stimuli or events in the outside world caused the neuron to fire action potentials. By so characterizing many neurons from a given area, one can determine what that area represents about the world.

To determine what stimuli cause the neuron to fire, the experimenter is typically guided only by his or her own intuition. For example, if one is exploring a piece of visual cortex, one knows that the neurons will be sensitive to visual stimuli, and so one may try many different types of visual stimuli to try to determine that which elicits the most spikes from the neuron. But this search is determined and limited by the experimenter's subjective sense of what is important. Alternatively, one can show a Gaussian noise ensemble (white or colored noise), and then principled methods exist to determine stimulus features to which the neuron is sensitive (by finding the spike-triggered average stimulus; and/or, by finding the eigenvectors of the spike-triggered covariance matrix with eigenvalues that are significantly perturbed from those of the unconditioned stimulus covariance matrix; in both cases, one must also correct for correlations in the stimulus ensemble). But these methods fail for stimuli other than Gaussian noise, such as an ensemble of scenes from a natural visual environment like that in which visual neurons evolved to function (I will call such an ensemble "natural scenes" for short).

These problems may not seem acute for studies of primary sensory structures. For example, in primary visual cortex, experimenters' intuitions have sufficed to characterize cells in important ways using simple stimuli such as luminance bars and gratings; and Gaussian noise can reasonably drive cells so that more principled methods can be used. Nonetheless, because cells adapt their response properties to stimulus statistics, it remains possible that a cell's responses to natural scenes are determined in significantly different ways than its responses to Gaussian noise or to simplified stimuli, and there has been no method to characterize such differences if they exist. Furthermore, as one goes further from the sensory periphery, at least in mammalian neocortex, the problems with existing methods become more and more acute: neuronal responses come to be more and more specialized and dependent on more and more complex features of the stimulus, and only very subjective methods currently exist to characterize this; and neurons have little or no response to unstructured noise, so that methods based on Gaussian noise stimuli cannot be used. For all of these reasons, one would like to be able to expose sensory neurons in an arbitrary area to natural scenes – a natural environment like that in which the neurons evolved, which should be able to drive the neurons regardless of the complexity of their response properties – and to determine in an objective and principled fashion those aspects of the stimulus that caused the neuron to respond.

It is this problem that Tatyana has tackled. She began with the very reasonable assumption that a neuron is sensitive to only a low-dimensional subspace in the high-dimensional space of possible stimuli – that is, to a relatively small set of "features", relative to the full set that defines all possible stimuli. An arbitrary pattern or feature – for the case of visual stimuli, this would be an arbitrary luminance pattern, for example an oriented pattern of light alternating with dark – represents a "dimension" in the high-dimensional space of all possible visual stimuli, and so we will refer to the small set of features or patterns to which a neuron is sensitive as the neuron's "relevant dimensions" – the dimensions in stimulus space that are relevant to determining its response. Tatyana then reasoned that the projections of the stimulus on the neuron's relevant dimensions will carry information about the neuron's response, whereas projections onto dimensions to which the neuron is insensitive will show no relationship to the neuron's response. That is, if one looks at the value, vs. time, of the projection of a given dimension onto the stimulus, and also looks at the neuron's response vs. time, these two time series will show some statistical dependence on one another for relevant dimensions, whereas they will be statistically indpendent for non-relevant dimensions. This statistical dependence can be characterized by the mutual information, as defined by Shannon, between the two time series – more mutual information in some sense indicates a greater dependence, while lack of mutual information indicates independence.

Tatyana therefore set out to find those stimulus directions whose projections onto the stimulus carry the most mutual information with the neuron's response. which she calls the method of "maximally informative dimensions". The idea, in sum, is that one will show an ensemble of stimuli, such as natural scenes for a visual neuron, and then extract from the neuron's responses the dimensions in stimulus space that were most informative about the neuron's response. Thus, for the first time, Tatyana's method allows one to stimulate a neuron with the stimuli it evolved to handle (or any stimulus ensemble), and to then extract in a principled way those features of the stimulus ensemble to which the neuron was sensitive.

Tatyana has brilliantly both formulated and implemented this approach, so that she is now using it to analyze our recordings from neurons in primary visual cortex. To do so, she computes the gradient of the mutual information in stimulus space, and uses simulated annealing to do a stochastic search up the gradient to find the most informative stimuli. The stimulus space is large, 900 pixels by 5 time slices, and so making the implementation work in practice is highly nontrivial. There are also difficult problems of overfitting – fitting chance rather than systematic correlations between the finite samples of stimulus and response – that Tatyana has carefully tackled. She has developed model problems on which to test her methods and to determine how their success scales with the amount of data available, and with reasonable amounts of data successfully finds the relevant directions for model neurons whose responses depend nonlinearly on one or two stimulus dimensions. She developed natural visual scene ensembles and worked on experiments to show these as stimuli while recording from cat visual cortex. From this data, she has so far been able to compute the single most informative dimension of neurons, and in some favorable cases where we have enough data the two most informative dimensions. With further work on experimental design, it should be possible to collect more spikes from individual neurons and thus to extract more dimensions. She has compared her results with characterizations of the same neurons by white noise or by drifting luminance gratings, and shown that the first stimulus dimension extracted by her method from natural scenes is similar to that extracted as the spike-triggered average (STA) to white noise and actually does somewhat better than the noise STA in predicting responses to gratings. Furthermore the most informative dimension to natural stimuli generally carries considerably more information about the neuron's response to natural stimuli than the noise STA carries about the neuron's response to noise. She has also compared the information about the neuron's response carried by the dimensions she finds to the total information about the stimulus carried by the neuron's single-spike responses; the latter can be calculated from the modulations of the neuron's firing rate in response to repeated presentations of the stimulus ensemble. For many cells, a single direction carries 30-60% of the total information carried by the neuron's action potentials, showing that she is indeed converging on highly informative stimulus dimensions.

While she is still in the process of analyzing the results, this method promises to have a major impact on the study of at least primary sensory areas of the cerebral cortex. In primary visual cortex, for example, neurons have been divided into two types, "simple" and "complex". Simple cells have responses that are roughly linear (up to rectification) in the luminance stimulus, that is, the responses of a simple cell to a general stimulus can be reasonably predicted by mapping out its responses to small spots of light or dark. A simple cell will typically respond best to a light bar of a preferred orientation at one location or to a dark bar of the same orientation at an adjacent, nonoverlapping location, that is, there is a spatial alternation between light-preferring and dark-preferring subregions. This pattern of light and dark preference can typically be mapped out by finding the STA to white noise. Complex cells, on the other hand, have responses that are in essence nonlinear, and typically show a good deal of position invariance: at most positions in the receptive field, the cell will respond to both a light bar and a dark bar of the preferred orientation. Until recently no systematic method of finding complex cell receptive fields has existed (the STA to noise typically has little or no structure, because at any given receptive field location either bright or dark bars can cause spikes, and so on average the stimulus before a spike is gray throughout the receptive field). Study of the spike-triggered covariance matrix in response to Gaussian noise has recently been introduced as one method of finding the dimensions to which a complex cell is sensitive, though results have thus far been limited to finding relevant patterns across one spatial dimension and time. Tatyana's method constitutes a second and more general method that is usable with any type of stimulus ensemble including natural scenes, and that allows extraction of patterns across two spatial dimensions and time with practical amounts of data. This method should allow us to characterize complex cells and their variety in a manner that has been heretofore impossible. The same method should work in primary sensory cortex of any modality, allowing similar leaps forward in somatosensory and auditory physiology.

A traditional method of distinguishing simple and complex cells is to measure the degree of temporal modulation of a neuron's response to a drifting grating – a simple cell shows a stronger temporal modulation over a grating cycle than a complex cell (in particular, one measures the F1/DC ratio of the neuron's response time course, where the F1 is the Fourier component at the same temporal frequency as the stimulus, and the DC is the mean response level; cells with ratios > 1 or < 1, are taken to be simple cells or complex cells, respectively). However, this is essentially a test of linearity, and simple cells are not truly linear. Tatyana proposes instead to distinguish simple from complex cells by the percent of information about the neuron's response carried by the single most informative dimension – the response of simple cells should largely depend on a single dimension, while complex cells must be sensitive to two or more dimensions to achieve position invariance. Indeed, preliminary results suggest that almost all cells for which a single direction carries more than 30% of the total information in the cell's response are simple cells by the traditional criteria, though both simple and complex cells by the traditional measure may have less information in their best direction. This suggests a new criterion for defining simple cells in a manner independent of the nonlinearities in the simple cell's responses.

"Higher", rather than primary, sensory cortical areas often have stronger invariances in their responses, *e.g.* a visual neuron in one area may respond to a particular face in almost any location or orientation. This would represent a sensitivity to many stimulus directions, the directions into which one stimulus is mapped under the group of invariances. Tatyana's method currently would probably not work in such areas. After completing her analysis of primary visual cortical neurons, Tatyana plans to tackle the problem of applying the method to areas that show such invariances. It is a difficult problem but she has both the brilliance and the tenacity to solve it. As already mentioned, in higher sensory areas, the problem of having no principled and objective way to determine what neurons are sensitive to has been a particularly critical one. Thus, if Tatyana succeeds in applying her method to such areas, it will be a major, major breakthrough in sensory physiology.

I have gone into quite a bit of detail as to the nature of Tatyana's research. I hope what comes across is the combination of very clear thinking about basic ideas, along with the theoretical, computational, and overall intellectual power to master a large set of complexities and difficulties and turn these ideas into reliable mathematical and computational implementations that shed new light on basic biological questions. Without a tremendous amount of intellectual firepower, focus, will power, and hard work, the many problems of developing practical solutions that give reliable and well-tested results would have been overwhelming. Without very clear thinking about the basic issues and the biology, either the original idea would have never been developed or the implementation would have become so bogged down in details that real biological results would never have been obtained. Very few if any of the postdocs that I have known over the years, including many selected by our Center with physics Ph.D.'s and outstanding records, would have been capable of doing what Tatyana has done.

Tatyana has recently applied for a K25 mentored career development award from the NIH for individuals moving from quantitive fields to biology. She wrote an outstanding, very clear and clearly reasoned proposal about her research, and received an outstanding score (153) suggesting that this award will be funded. This award will provide salary funding and some research support through her remaining years of mentored research, and augurs very well for her ability to obtain research funding in the future.

Tatyana is a wonderfully clear thinker and speaker. She has a complete mastery of the material she studies and speaks with deep insight and complete clarity. I believe she will be an outstanding teacher. (Her English, both written and spoken, is excellent, except for some small problems with knowing when to use or omit articles – "the", "a", and the like.) She works incredibly hard, often the first here in the morning and the last to leave at night, and with amazing focus. She is quite simply an outstanding and focused theoretical physicist who has creatively brought her skills into a new field with outstanding results. She is also humble and soft-spoken and a joy to work with. Only two years after she completed her physics Ph.D., her work in her new field is already having a significant impact on brain studies. I fully expect her to have a large and growing impact in the future and to continue to make new and major contributions.

Tatyana is an outstanding and brilliant theorist, one of the best scientists I have ever known. It is still early days in her neuroscience career, with few publications yet (though a major publication on her analysis of visual cortical cells should be submitted in the next few months), and thus far she has worked on only one major problem in her new field. As a result, whoever invests in her now will be getting someone absolutely outstanding well before her real market value is fully apparent. I recommend her to you as strongly as possible.

Sincerely,

Ken Miller

Kenneth Miller