Biocomplexity Faculty Search Committee
c/o Prof. Rob de Ruyter van Steveninck
Biocomplexity Institute
Indiana University
Swain Hall West 117
Bloomington, IN 47405-7105

December 22, 2003

Dear Prof. Rob de Ruyter van Steveninck:

I'm sending you my curriculum vitae and a statement of my research interests, pursuant to applying for a tenure-track assistant professor in Biocomplexity. I have also arranged three letters of recommendation to be sent to you directly.

I obtained a MS in theoretical physics in 1995, and Ph.D. in electrical engineering in 1997. I worked with Professor David Mumford for a year as a postdoctoral fellow at Brown on statistics of natural scenes and its relation to vision, a topic that I have continued to work on. After I left Brown in 1998, I moved to University of Arizona to work with Professor Richard Zemel on probabilistic models of visual cue combination. Since September 1999, I have been working with Professor Dale Purves here at Duke University on a probabilistic framework on visual perception and visual system structure and function. By virtue of this interdisciplinary training at undergraduate, graduate and postdoctoral levels, I now have a very solid background in mathematics, physics, computational modeling, cognitive science and neurobiology, and have published more than 40 papers in peer reviewed journals and conferences on topics such as brightness perception, motion perception, stereo vision, space perception, shape perception, perceptual organization and statistics of natural scenes.

I'm especially interested in this position. My research goal is to combine advanced natural scene analysis, computational modeling with psychophysical experiments, imaging and electrophysiological recording to thoroughly explore a wholly probabilistic framework for understanding visual brain function. This work will entail exploring the connections and organization of the visual cortex, the dynamic patterns of neural activity elicited by visual stimuli, the computations performed by the visual system and perceptions generated thereby.

I look forward to hearing from you.

Yours sincerely,

Zhiyong Yang

Zhiyong Yang, Ph. D.
Center for Cognitive Neuroscience
LSRC Building, Box 90999, Duke University
Durham, NC 27708, USA
Phone: 919-684-6276; Fax: 919-681-0815
Email: zhyyang@duke.edu

Statement of Research and Teaching Interests

Zhiyong Yang, Ph.D.
Center for Cognitive Neuroscience
Box 90999, LSRC Building, Duke University
Durham, NC 27708-0999, USA
Phone: (919) 684-6276, Fax: (919) 681-0815
Email: zhyyang@duke.edu

My long-term goal is to develop a theoretical framework of cortical function in primates (using vision as a model), and to apply this framework to advancing both computer technology and human-machine interactions.

Ever since Hubel and Wiesel's groundbreaking work on neuronal receptive-field properties, the function of the visual brain has been taken to be the extraction and analysis of the features of a scene. This conception of neuronal properties has become increasingly difficult to square with a fast growing body of pertinent neurobiological data. With respect to computation, many studies of the principles underlying visual perception continue to be placed in the framework developed by David Marr some 20 years ago, even though Marr's paradigm has not succeeded in rationalizing the relationship between the response properties of visual neurons, visual system structure or visual perception. As a result of these frustrations, there has been a resurgence of interest in the relation of the statistics of natural environment and neuronal properties, based on the assumption that sensory systems have evolved to efficiently code features in natural environments. With a few exceptions, however, significant insights into visual system structure, function and the strategies that could generate percepts have not yet been forthcoming from this work. In short, several avenues of mainstream vision research are, by the admission of many workers in the field, at an impasse.

The broad hypothesis that I intend to pursue is that the response properties of visual neurons and their connections, the organization of visual cortex, the patterns of activity elicited by visual stimuli and visual perception are all determined by the probability distributions of the possible physical sources underlying visual stimuli. To fully explore this admittedly challenging idea, I intend to pursue several research paths concurrently:

1) *Develop probabilistic models of neuronal response properties, neuronal connections, cortical organization and patterns of activity elicited in the visual cortex using advanced computational modeling.*

Since observers can access the visual world only through images projected onto the retina, the physical sources of visual stimuli, on which successful behavior depends, cannot be known directly. This problem could be solved by incorporating into the visual system a probabilistic strategy to generate visual perception based on the probability distributions of the sources underlying visual stimuli. In this conception of vision, neurons are not detecting or encoding photometric features, but by virtue of their activity levels are acting as estimators of the probability density of the possible physical sources underlying any given stimulus. From this perspective, the function of visual cortical circuitry is to propagate, transform and combine probability distributions. The iterated structure of primary visual cortex in primates may thus be organized in the way it is in order to generate probability distributions pertinent to the simpler aspects of visual stimuli. By the same token, extrastriate visual cortical areas may serve to generate joint probability distributions of the physical sources underlying visual stimuli by propagating and combining the probability distributions elaborated in V1. The activity patterns elicited by any visual stimulus would, in this conception, be determined by the joint probability distributions of physical sources underlying the relevant aspects of visual stimuli and the underlying operational dynamics.

To explore this concept of visual cortical function, I will therefore develop probabilistic models for these processes. In particular, I will examine how neuronal properties in response to orientation, color, motion and stereoscopic disparity are related to the probability distributions of the physical sources, asking whether contextual effects on these neuronal responses can be understood in terms of their influence on the relevant probability distributions. I will also examine whether cortical connections and their anatomical organization (i.e., orientation maps, ocular dominance columns, blobs) can be rationalized in these terms.

2) *Test these probabilistic models using electrophysiological recording and imaging in experimental animals.*

To test these probabilistic models, I will design a set of stimulus ensembles for which the probability distributions of the underlying physical sources are known (some prior probabilities and conditional probability distributions needed for this purpose will be obtained from the statistical analysis of natural scenes described below). In each ensemble one or more physical variables will be systematically manipulated. These stimuli will then be presented to experimental animals and the neuronal responses, in collaboration with other colleagues, will be recorded using electrophysiological and/or optical imaging technology. Advanced statistical learning theory will then be used to examine whether the neuronal responses reflect the relevant probability distributions. More specifically, this will be done by examining: 1) whether the probability distributions obtained from neuronal activity tallies with the known probability distributions underlying the stimuli; and 2) whether the neuronal activity generated from known probability distributions tallies with the neuronal activities actually recorded.

3) *Study the statistics and probabilistic models of natural scenes using range, hyper-spectral, and stereoscopic information derived from natural visual environments; examine the relation of the analyses of these databases to visual system structure and function, and to visual perception.*

I plan to use a hyperspectral imaging system, laser range scanning and a stereoscopic imaging system (these systems have only recently become available and affordable) to acquire a large database of images of natural scenes from indoor, outdoor and entirely natural environments. I will also use a spectroradiometer to measure illuminant spectra at distributed locations in each scene. Thus each image in this unique database will provide related information about the spectral return for each wavelength band, as well as the range and disparity at each pixel in the images. The spectrum of the illuminant in the scene will provide the needed information about the lighting. Another database of natural image sequences will be acquired using a digital video camera to study motion perception.

I will then take advantage of the techniques of dimension reduction, probability density estimation and graphical probabilistic modeling to develop a framework for computing the probability distribution of the sources of natural images, using a minimum of assumptions about the properties of the objects or conditions in the relevant scenes. I will subsequently test whether the relevant aspects of perception of color, brightness, space, stereoscopic depth and motion (e.g., the structure of color space and visual space) can be accounted for by the statistics of image-source relationships in natural visual environments. I will also examine whether and how the relevant aspects of neuronal properties and cortical organization (e.g., modulation of neuronal response properties, contextual effects, long-range connections) are shaped by the statistics of natural environments.

4) *Study the perception of brightness, color, motion, stereoscopic depth, shape, space, objects and scenes under realistic conditions using both advanced probabilistic modeling and psychophysical experiments.*

For each of these domains of visual perception, I will create a set of stimulus ensembles. In each ensemble, one or more physical variables will be systematically manipulated. To compute the probability distribution of physical sources underlying the stimuli, I will use graphs to represent the probabilistic relations between physical variables. The advantages of using graphical techniques are: 1) the probability relations between variables in space-time are made clear; 2) techniques used in data mining, knowledge engineering, and expert systems can be applied to the problem; and, 3) by mapping back the operations on the graphs to the stimuli, psychophysical experiments can be designed to test the detailed predictions of the theory. The prior probabilities and conditional probability distributions needed for this purpose will be obtained from the statistical analysis of natural scenes, as already described. I will then test in human subjects whether visual perception in each of these domains is indeed generated on basis of the probability distributions of physical sources underlying visual stimuli.

5) Study visual perception and visual system structure and function by examining digital lives evolving in realistic natural environments.

I will design autonomous digital agents and let them evolve in an unsupervised manner in realistic natural environments (i.e., the database of natural environments I will have acquired). I will then examine whether the evolved virtual organisms develop visual capabilities similar to those evident in human perception, and if so how this is accomplished by their system architecture and operational dynamics.

6) *Examine whether patterns of visual cortical activity elicited by visual stimuli monitored by fMRI in humans (and/or optical imaging of fMRI in experimental animals) and the corresponding percepts can be understood in terms of the relevant probability distributions of physical sources underlying visual stimuli.*

I will examine patterns of visual cortical activity monitored by fMRI in humans (and/or experimental animals) in response to: 1) simple geometrical stimuli that elicit anomalous percepts (i.e., geometrical illusions); 2) simple luminance/color stimuli that elicit lightness contrast/constancy or color contrast/constancy effects; 3) simple stereoscopic stimuli that generate percepts difficult to account for in terms of retinal disparity; 4) simple moving stimuli that generate percepts profoundly influenced by various contextual cues. The question I will ask is whether these monitored responses can also be understood in terms of the relevant probability distributions of physical sources underlying visual stimuli. Because these stimuli will be relatively simple and parametrically controlled, the relevant probability distributions can be easily computed. I will then use advanced probabilistic modeling to examine the relation between these probability distributions and the corresponding brain activity.

Importantly, these experiments can be generalized to a variety of more natural, complex stimuli. This set of stimuli will be designed such that the complex stimuli contain components that correspond to a set of simple stimuli such as those described above. For both simple and complex stimuli, I will compute the probability distributions of underlying physical sources (as before, some prior probabilities and conditional probability distributions will be obtained from the statistical analysis of natural scenes described above). Statistical computing will then be used to examine whether the cortical activity observed in response to simple stimuli is predicted by the underlying probability distributions, and whether the cortical activity observed in response to complex stimuli is predicted by combining the probability distributions of the relevant components. Performing this latter analysis will entail two related procedures: 1) examining how the probability distributions are represented in the activity of a population of neurons; and 2) examining how the operations on probability distributions are mapped onto the operations that generate activity patterns elicited by complex stimuli in terms of the patterns elicited by simple stimuli.

7) *Develop biologically inspired prototype visual pattern analysis and recognition systems*

Finally, I will try to develop prototype systems for visual pattern analysis and recognition based on this probabilistic framework for understanding of the visual neuronal properties, cortical connections and the organization of visual cortex. Such a system should be far superior to existing vision systems, and would have the potential to provide advanced pattern analysis and recognition systems for a variety of applications, including biomedical image (MRI) analysis.

In conclusion, I plan to combine an extensive statistical analysis of natural environments, advanced probabilistic modeling with animal studies, brain imaging and psychophysics to thoroughly explore a novel framework for understanding the visual system structure, the computations it performs and perceptions generated thereby. If this is indeed the framework in which vision must be understood, then the interdisciplinary work I intend to pursue should go far towards establishing a novel paradigm for understanding and interpreting the vast amount of information about primate visual system structure and function that has been obtained in recent decades, but which has so far been difficult to relate to what humans and other primates actually see. More generally, if this statistical strategy is the way other sensory modalities work, the framework I have summarized here should dramatically affect present understanding of brain development, organization, and function.

With respect to my teaching interests and abilities, I would like to teach at both the undergraduate level (introduction to statistical methods, introduction to perception and sensation, neural computation, neuroscience and cognitive neuroscience) and graduate level (information theory, signal processing, image processing, computer vision, pattern recognition, neural computation, advanced statistical methods, Bayesian statistics, graphic probabilistic models, pattern theory, theoretical neuroscience and vision science). I would also like to lead seminars on some advanced topics of vision science (neurobiology of vision, computational models, perceptual organization, natural scene statistics, probabilistic models of cortical structure and function) and other topics as these mature on the basis of my own and others' research. Finally, I'm fully committed to supervising undergraduate, graduate and postdoctoral students.