Yan Karklin

Research associate
Howard Hughes Medical Institute
New York University

4 Washington Pl, Rm 809
New York, NY 10003
212 992-8751
yan.karklin@nyu.edu

citizenship: USA

Education:

Ph.D., Computer Science Department & Center for the Neural Basis of Cognition 2000-2007
Thesis: *Hierarchical statistical models of computation in the visual cortex*, Advisor: Michael S. Lewicki
Carnegie Mellon University, Pittsburgh, PA.

B.A., Neuroscience, *magna cum laude* 1996-2000
Thesis: *Spike timing precision in frog motoneurons,* Advisor: Stephen A. George
Amherst College, Amherst, MA

Academic positions:

Research associate Feb 2008-present
Howard Hughes Medical Institute & Center for Neural Science, Supervisor: Eero P. Simoncelli
New York University, New York, NY

Visiting researcher May-Aug 2003
Physical Biosciences Division, Supervisor: Stephen R. Holbrook
Lawrence Berkeley National Lab, Berkeley, CA

Visiting researcher Jan -Jul 1999
Dept. of Neurobiology, Supervisor: Ehud Ahissar
Weizmann Institute of Science, Rehovot, Israel

Visiting researcher May-Aug 1998
Dept. of Stomatology, Medical Center, Supervisor: Peter Sargent
University of California, San Francisco, San Francisco, CA

Honors and Awards:

DOE Computational Science Graduate Research Fellowship 2002-2006

NASA Graduate Researchers Program Fellowship 2002 (declined)

Sigma Xi Research Society Membership 2000

National Merit Scholarship 1996

Teaching Experience:

Guest Lecturer, Computational Perception and Scene Analysis — Spring 2004, 2006

Teaching Assistant, Algorithm Design and Analysis — Fall 2001

Teaching Assistant, Computer Vision — Spring 2002

PUBLICATION LIST

Refereed journal articles:

Y. Karklin and E. P. Simoncelli, Efficient coding with biological constraints explains organization of retinal ganglion cells, *in preparation*.

Y. Karklin and E. P. Simoncelli, Optimal transfer function in a noisy nonlinear neuron, *in preparation*.

Y. Karklin and M. S. Lewicki, Covariance factor model for continuous data, *in preparation*.

Y. Karklin and M. S. Lewicki, Emergence of complex cell properties by learning to generalize in natural scenes, *Nature*, 457, 83-86, 2009.

Y. Karklin and M. S. Lewicki, A hierarchical Bayesian model for learning non-linear statistical regularities in non-stationary natural signals, *Neural Computation*, 17 (2): 397-423, 2005.

Y. Karklin and M. S. Lewicki, Learning higher-order structures in natural images, *Network: Computation in Neural Systems*, 14: 483-499, 2003.

Refereed conference proceedings and abstracts:

Y. Karklin and E. P. Simoncelli, Efficient coding of natural images and movies with populations of noisy nonlinear neurons, *Computational and Systems Neuroscience (CoSYnE)*, 2012.

Y. Karklin and E. P. Simoncelli, Efficient coding of natural images wit ha population of noisy linear-nonlinear neurons, *Adv in Neural Information Processing Systems (NIPS)*, 2011.

Y. Karklin and E. P. Simoncelli, Optimal information transfer in a noisy nonlinear neuron, *Computational and Systems Neuroscience (CoSYnE)*, 2011.

Y. Karklin and M. S. Lewicki, Is early vision optimized for extracting higher-order dependencies?, *Adv in Neural Information Processing Systems (NIPS),* 2006.

Y. Karklin, R.F. Meraz, and S. R. Holbrook, Classification of non-coding RNA using graph representations of secondary structure. *Proc of the Pacific Symposium on Biocomputing 10*, 2005.

Y. Karklin and M. S. Lewicki, A model for learning variance components of natural images, *Adv in Neural Information Processing Systems (NIPS)*, 2003.

Invited reviews:

Neural Information Processing Systems

Neural Computation

Journal of Vision

Journal of Computational and Graphical Statistics

Scientific courses:

Cold Spring Harbor Laboratory, Computational Neuroscience: Vision, Jun 2008

Invited talks and presentations:

NCL, Albert Einstein College of Medicine, Mar 2012

Gordon Research Conference, Bates College, Aug 2010

Fellows seminar, NYU, Dec 2010

ICML Workshop on learning feature hierarchies, Jun 2009

Psychology Department, NYU Jan 2009

Computer Science Department, NYU, Dec 2008

CIFAR workshop, May 2008

Center for Neural Science, NYU, Jul 2007

Gatsby Computational Unit, UCL, Jul 2007

Redwood Institute, UC Berkeley, Jun 2007

Student seminar series, Computer Science Dept, CMU, May 2007

CNBC Brain Bag series, CMU, Feb 2007

Department of Energy Computational Science Fellowship Conference, July 2006

Center for the Neural Basis of Cognition Retreat, CMU, Oct 2006

Department of Energy Computational Science Fellowship Conference, July 2005

Department of Energy Computational Science Fellowship Conference, July 2004

Department of Energy Computational Science Fellowship Conference, July 2004

Yan Karklin, PhD
RESEARCH STATEMENT

I study sensory processing of complex natural scenes. I am especially interested in representations at the early and mid-level sensory stages – how are elements of the environment encoded in the periphery, and how is this information transformed into representations of meaningful scene properties, ones that can ultimately guide behavior? My work has mainly focused on vision, with some application to audition, but I am also keen to explore other domains.

Natural scenes are richly structured, and at the same time tremendously complex and varied. Our experience of the world relies on signals gathered at the sensory organs, but by the time light arrives at the eye and sound waves reach the ear, relevant pieces of information are deeply entangled and buried in noise. Hiking through the woods, we are startled and look back; a flicker of light on the retina, a deflection of the basilar membrane in the cochlea – is that our companion's brown jacket and heavy breathing, or something wild and threatening? In order to deal with this uncertainty and draw inferences about the state of the world, sensory systems must complement the incoming stream with prior knowledge of the surroundings, so it is crucial that computational strategies are adapted to the specific environment (bears do not wear plaid, so we are safe.) It is well known that the brain solves such perceptual tasks far better than any machine. In fact, it does so under severe constraints. Neural processing consumes metabolic resources and relies on an architecture composed of units that are not perfectly reliable. Understanding the structure of sensory signals, as well as the constraints relevant to biological computation, allows us to formulate and test concise theories of function. These are the recurring themes of my research: the complex regularity of natural scenes, statistical inference under conditions of noise and uncertainty, and the interplay between optimal behavior and biological constraints.

As part of my thesis research under the mentorship of Michael Lewicki, I developed a functional theory of computation in complex cells in primary visual cortex (V1). Previous work had showed that basic properties of simple cells in V1, such as orientation and frequency selectivity, could be explained using the principle of efficient coding and a linear model adapted to natural image statistics. However, it has been much more difficult to extend this theory to neurons that do not code simple linear relationships in their inputs. Motivated by the observation that perceptually distinct image regions (e.g., textures, object boundaries) are better discriminated using local statistics than precise pixel values, we proposed that complex cells subserve invariant image representation by encoding these statistical relationships, and constructed a computational model based on this theory. This description of complex cells is very different from the standard energy model, in that it is based on an explicit computational goal (representational invariance) and it incorporates crucial elements of uncertainty, noise, and statistical inference. Training the model on natural images, we found that, without making any assumptions about neurons or their properties, we could predict a number of known non-linear neural behaviors, including phase invariance, cross-orientation suppression, and orientation-dependent surround suppression. Because these behaviors emerge as optimal solutions for achieving invariance, our model was able to provide a necessary functional explanation for these disparate effects.

The model we developed also makes novel, testable predictions for how neural responses are modulated by context in natural images, a setting in which traditional accounts of surround effects (derived using simple stimuli like gratings) typically fail. I am beginning to carefully test these predictions by analyzing adaptive normalization behavior in V1 neurons, in data collected by Ruben Coen Cagli. In this experiment, neurons are presented with images either in isolation, or embedded in their natural background, and modulations in their response are related to image properties inferred by the model, rather than image pixels. This integration of theory and experiment promises to advance our understanding of neural coding of natural images and guide further refinement of the computational model.

Another focus of my recent work has been information processing in the retina. Decades of anatomical and physiological studies have painted an exquisitely detailed picture of the primate retina, yet a functional understanding of the most basic properties is still missing. Why are the receptive fields of retinal ganglion cells localized and unoriented, with a center-surround opponent sensitivity? What roles do ganglion cell subtypes such as On and Off-center cells, play in conveying visual information to the rest of the brain? And what about the seemingly redundant mosaics of receptive fields of different ganglion cell types that independently tile the visual space? Together with Eero Simoncelli, I showed that by extending an efficient coding model to include neural noise, output nonlinearity, and limited metabolic resources, we can derive many of the basic properties of retinal organization: center-surround opponency, rectifying output nonlinearities, tiling On- and Off-center receptive fields, and even asymmetries between the On- and the Off- populations that have been observed in experiments. As in the model of complex cells, these predictions depended on the statistics of natural images and neural representation of uncertainty.

These results open up three directions of inquiry. First, I intend to use this framework to gain insight into other nonlinear transformations in the retina, such as contrast gain control and adaptation to stimulus statistics. Adaptive behaviors have been characterized in elegant *in vitro* experiments, but the theoretical framework I have developed allows us to answer several specific questions. For example, what is the optimal extent of spatial integration over which a neuron is to compute local contrast to adjust its gain? Should the On- and Off-center cells pool this information? And what components of the circuit (e.g., bipolar or amacrine cells) subserve this computation? I have already made some progress in addressing these questions (and obtained exciting initial results), but I expect this pursuit to guide a long term investigation.

The second natural extension of this work is to incorporate temporal processing into our analysis. This is not trivial. It is technically challenging to estimate information conveyed by a population of neurons over time, and even settling on "natural" spatiotemporal data involves difficult modeling choices. Nevertheless, I am confident that we can make headway, and I believe this work will shed light on a major mystery of early sensory processing: the dual parvo- and magnocellular processing streams. (Their temporal properties differ significantly, as do the nonlinear and adaptive behaviors.) Theoretical advances here will also have direct bearing on coding strategies in the auditory pathway. I have begun working on a spatiotemporal model, using temporal filters on a coarse scale. Even this crude model makes interesting predictions, such as increased latency in the spatial surround of the receptive field relative to its center.

The third line of research, about which I am most excited, aims to integrate these advances into a multistage hierarchical model encompassing the visual cortex. A major limitation of current models of cortical processing is that they ignore computation carried out in early sensory stages. When modeling V1 (or, for that matter, building feature hierarchies for computer vision systems), responses are computed directly on stimulus pixels, bypassing the retina and everything else on the way. By working on the outputs of a more realistic retinal model, such as the one described above, we can hone in on the properties (and the constraints) specific to cortical computation.

In addition to providing insight into neural computation, my work has lead to the development of sophisticated statistical models of complex, high-dimensional data, with obvious applications in image and sound analysis, denoising, segmentation, and object recognition. Recent advances in the computer vision community (e.g., using local normalization to improve object recognition performance) parallel my planned research into retinal mechanisms of gain control. Thus, I am very interested in bridging these fields, bringing approaches that prove useful in neuroscience into the domains of signal processing and computer vision, and vice versa. In this respect, as in others, Brown would be a wonderful environment, with its deep strengths in computer science, applied math, and visual and systems neuroscience.

Yan Karklin, PhD
TEACHING STATEMENT

I am enthusiastic about teaching and look forward to engaging with students, giving lectures, and developing new courses. Some of my best learning moments in graduate school occurred in lectures delivered by passionate and skilled instructors, when a concept suddenly became clear and familiar, and I appreciate the value of a prepared and an attentive teacher. Other times, an effective course simply provided a solid overview of the field or a set of tools needed for my research. The process of teaching is beneficial for the instructor as well, providing on opportunity to solidify the grasp of wider topics, or to inspire new research directions. In general, I find interaction with students and junior researchers stimulating – discussions often lead to new insights, and sometimes to new research projects and fruitful collaboration with students and other faculty.

I have acted as Teaching Assistant (teaching sections, developing homework, delivering guest lectures, and grading assignments) in two graduate level courses, Computer Vision, and Algorithms. The course on Algorithms at Carnegie Mellon was a relatively high level course, and quite demanding for someone with a neuroscience background. It was a challenge, but there was no better way to immerse so quickly and so thoroughly in a new topic. The course in computer vision was more familiar territory, so I was able to focus on tailoring the presentation of course material and the design of assignments to student needs. In graduate school, I also delivered guests lectures and assisted with the preparation of a graduate level course on computational scene perception.

I would be comfortable teaching most undergraduate courses in neuroscience and computer science, as well as graduate level material on visual, systems, and computational neuroscience, visual perception, statistical methods, signal processing, machine learning, and computer vision. I am particularly keen to develop higher level courses along the lines of my interests and expertise, covering topics in neural computation, information theory and Bayesian methods, computational scene analysis, and image processing.