Xaq Pitkow
Postdoctoral Fellow, Department of Brain and Cognitive Science
Meliora Hall, University of Rochester, Rochester NY 14627
Tel 978-460-5144 • xaq@post.harvard.edu

Faculty Search Committee
Department of Neuroscience
Brown University
Providence, RI

April 25, 2012

Dear Members of the Search Committee,

I am writing to apply for your faculty position in Computational Neuroscience. I am currently a postdoctoral researcher in the Department of Brain and Cognitive Science at the University of Rochester, where I study neural mechanisms of perceptual inference.

My fundamental goal is to understand how the brain perceives the external world, based only on our inherently ambiguous sense data. I am drawn to the hypothesis that we are constantly performing probabilistic inference, because it has substantial experimental support across sensory, cognitive, and motor domains, and an appealing power for generating new predictions. But connecting this abstract notion to concrete neural mechanisms for specific tasks leaves much work to be done, so this is where I have been focusing my efforts. In particular, I study how approximate probabilistic inference can account for the neural transformation of low-level features into more abstract, higher-level object properties, in human, animal, and mathematically modeled brains.

My attraction to probabilistic inference grew out of my longstanding interest in neuroscientific explanations based on first principles. As my first project in Markus Meister's lab at Harvard, collaborating also with Haim Sompolinsky, I analyzed the optimal way to decode visual information blurred by our incessant eye jitter, and found that the mathematically ideal solution could be implemented by primary visual cortex. In my second project, I falsified the important theory that retinal center-surround receptive fields efficiently removed correlations in natural images, and then rehabilitated the theory in a more general form to show that the curiously sparse retinal responses were actually optimal for transmitting information efficiently. As a Swartz fellow at the Columbia Center for Theoretical Neuroscience, advised by Ken Miller and Larry Abbott, I began a multi-component project to relate object perception to statistical inference, and demonstrated that lateral connectivity in visual cortex improves the inference of object boundaries. When my wife took a faculty position at the Eastman School of Music, I continued this work at the University of Rochester, where with Alex Pouget I have been implementing these neural computations with population codes. Over the next few years I intend to apply these ideas to other sensory tasks, and to address the computational roles of both lateral and feedback interactions in probabilistic inference. Although my research has been primarily theoretical, clear predictions should be tested. During my graduate studies at Harvard I did electrophysiological and psychophysical experiments myself. As a postdoc I have established collaborations with two psychophysics labs (Michele Rucci at BU, Daphne Bavalier at UR) and two electrophysiology labs (Tony Movshon at NYU, Greg DeAngelis at UR) in order to test key theoretical predictions and develop models to explain data.

Brown University's Department of Neuroscience and Institute for Brain Science would be a fantastic place to conduct this research. I expect productive interactions with the many creative faculty studying neural function, certainly in vision and computation, but also more broadly in audition, olfaction, and motor control. I am extremely enthusiastic about the computational neuroscience community already established at Brown, and would energetically participate in interdisciplinary collaborations.

As a speaker I excel in making complex ideas comprehensible, and as a discussion leader I am good at identifying key questions, focusing debate, and fostering critical thinking skills. My students, colleagues and mentors attest that these attributes make me an inspiring and highly effective teacher. Three of my students have even credited me for inspiring them to go into neuroscience. Others have routinely sought me out for brainstorming and critical comments. One colleague actually professed that when he prepares a talk, he asks himself, "What would Xaq ask?" At Harvard I was a teaching assistant for classes ranging from applied math to molecular biology. I would enjoy mentoring graduate students and teaching courses in computational and systems neuroscience, as well as interdisciplinary courses in complex systems, machine learning, or statistical methods in the brain sciences.

Please find attached my CV, research and teaching statements, and a few reprints. References will be submitted by my current advisor, Alex Pouget; my postdoctoral advisors at the Center for Theoretical Neuroscience, Ken Miller and Larry Abbott; and my Ph.D. advisor Markus Meister. Please let me know if I can provide you with any other information. I look forward to hearing from you.

Sincerely,

Xaq Pitkow

xaq pitkow

Department of Brain and Cognitive Sciences
Meliora Hall, University of Rochester, Rochester, NY 14627
(978) 460-5144 • xaq@post.harvard.edu • columbia.edu/~zsp2101

POSITION

University of Rochester 9/2010–
Postdoctoral Research Scientist (advisor: Alexandre Pouget)

EDUCATION

Columbia University 1/2007–8/2010
Postdoctoral Research Fellow, Center for Theoretical Neuroscience
(advisors: Ken Miller, Larry Abbott)

Harvard University 9/1999–3/2006
Ph.D. in Biophysics
thesis: *Optimality Principles for the Visual Code* (advisor: Markus Meister)

Princeton University 9/1993–6/1997
A.B. in Physics, magna cum laude
thesis: *how to simulate Quantum physics on a Quantum computer using a Quantum lattice gas on a nonuniformly Quantized space* (advisor: Washington Taylor IV)

PUBLICATIONS

Beck J, Ma WJ, **Pitkow X**, Latham P, Pouget A (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74(1): 30–9.

Pitkow X, Meister M (2012). Decorrelation and efficient coding in retinal ganglion cells. *Nature Neuroscience.* 15(4): 628–35.

Pitkow X, Ahmadian Y, Miller KD (2011). Learning unbelievable probabilities. *Advances in Neural Information Processing Systems.*

Pitkow X (2010). Exact feature probabilities in images with occlusion. *Journal of Vision* 10(14): 42.

Pitkow X, Sompolinsky H, Meister M (2007). Visual acuity in the presence of fixational eye movements. *PLoS Biology*, 5(12): e331.

HONORS

Sloan-Swartz Fellowship in Theoretical Neurobiology	2007–2009
National Science Foundation Fellowship in Biophysics	1998–2001
Allen G. Shenstone Prize (Princeton Physics Thesis Prize)	1997

RECENT POSTERS

Pitkow X (2012). Implicit representation of high-dimensional stimulus distributions. *Computational and Systems Neuroscience.*

Pitkow X, Ahmadian Y, Miller KD (2011). Approximate inference by neural circuits. *Computational and Systems Neuroscience*.

Pitkow X, Ahmadian Y, Miller KD (2010). The value of lateral connectivity in visual cortex for interpreting natural images. *Computational and Systems Neuroscience*.

Pitkow X (2009). Inference of object attributes from local image features caused by occlusion. *Computational and Systems Neuroscience.*

Pfau D, **Pitkow X**, Paninski L (2009). A Bayesian method to predict the optimal diffusion constant in random fixational eye movements. *Computational and Systems Neuroscience.*

Pitkow X, Miller KD (2008). Calculation of arbitrary statistics of naturalistic scenes. *Computational and Systems Neuroscience.*

Pitkow X, Miller KD (2007). A Bayesian model of color fill-in. *Society for Neuroscience*.

REFEREEING

Neuron, PLoS Computational Biology, Frontiers in Computational Neuroscience, Neurocomputing, Journal of Neurophysiology, Cerebral Cortex, Computational and Systems Neuroscience, Neural Information Processing Systems

TEACHING

Guest lecturer, University of Rochester
- Intro to Computational Neuroscience (instructor: Alex Pouget) — 2011

Organizer, University of Rochester
- Computational Neuroscience reading group — 2011–

Guest lecturer, Columbia University
- Computational Methods in Neuroscience — 2008–2010

Teaching Fellow, Harvard University
- Computational Neuroscience (instructor: Haim Sompolinsky) — 2003
- Function of Neural Circuits (instructor: Markus Meister) — 2002
- Introduction to Molecular Biology (instructor: various) — 2001
- Applied Math: Complex Analysis (instructor: Efthimios Kaxiras) — 2000

OUTREACH

Mott Hall 7th grade mentor, Harlem, New York City	2007–2009
School of the Museum of Fine Arts, Boston	2006
University of Texas at Austin, Art Department	2007–2011
Art Institute of Austin	2009–2010

Statement of Research Interests

xaq pitkow

General approach

A wing would be a most mystifying structure if one did not know that birds flew. – *H.B.Barlow (1961)*

My scientific goal is to explain functions of neural circuitry based on computational principles. On the most general level, I explore the hypothesis that the brain evolved to perform optimal inference on behaviorally relevant perceptual tasks. Because sensory evidence is ambiguous, and properties of the external world may manifest as many different sensory patterns, the brain must use a statistical approach to perception: It must exploit learned prior knowledge about the world and its own sensory representations in order to select the most probable interpretations from a multitude of improbable ones. My research program is to mathematize perceptual tasks, identify statistical inference algorithms that perform comparably to animals, model these algorithms with neural circuits subject to biological constraints, and deduce key predictions about behavior and neural mechanisms. These predictions comprise distinctive properties of behavior or circuit function that are essential for the computational model to work correctly. **This 'normative' approach bridges different levels of explanation by connecting the computational to the mechanistic.**

To evaluate these general ideas, I have focused mostly on vision, and in particular on how the brain transforms low-level visual features like oriented edges into higher-level causal properties like object boundaries. I am also interested in other sensory systems, how neural representations are learned and transformed, and whether those representations are efficient in a quantifiable sense.

Past research

ACUITY IN THE PRESENCE OF FIXATIONAL EYE MOVEMENTS. In my thesis work, my collaborators Markus Meister, Haim Sompolinsky and I proposed a solution to a long-standing puzzle in vision science: How can we see such fine detail despite the fact that our eyes are moving all the time? These eye movements smear the image on the retina by tens of arc minutes, yet we can distinguish visual detail about 30 times finer. We proposed a solution to this mystery in the form of a probabilistic inference algorithm that can process neural signals coming from the retina and undo the smear. In essence this mechanism estimates from moment to moment where the image has shifted on the retina, and aligns the neural signals accordingly. This abstract algorithm can be readily implemented by biological circuits (Figure 1), and its characteristics accord well with many known properties of early visual processing. Psychophysical experiments I conducted showed that human performance on simple discrimination tasks were consistent with the performance of our model. Thus, we were led to propose that primary visual cortex (V1) solves the problem of blur from fixational eye movements.[1]

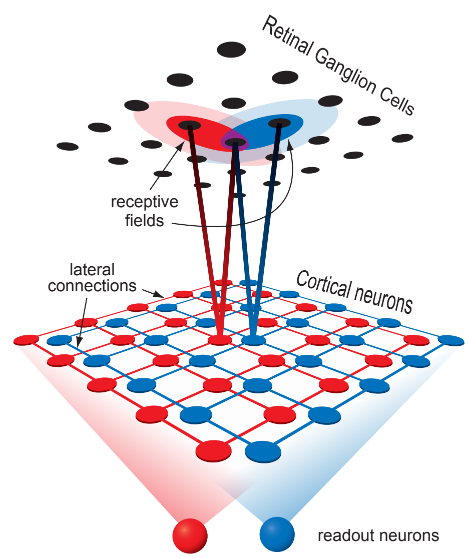

Figure 1. Neural network implementation of the optimal decoder for fixational eye movements. Spiking responses from retinal ganglion cells (black circles) provide inputs to orientation-selective neurons in cortex (red and blue circles). Cortical cells of a given orientation are mutually excitatory with a coupling strength related to the mean speed of fixational eye jitter. Their responses indicate the most probable positions and orientations given the stimulus history. Optimal readout of orientation is accomplished by spatial pooling (bottom).

DECORRELATION AND EFFICIENT CODING IN THE RETINA. 'Efficient coding theory' claims that sensory systems have evolved to transmit information as efficiently as possible within the laws of physics. This idea is one of the most successful theories in neuroscience, having derived from first principles explanations of phenomena as diverse as human contrast sensitivity and the sizes of insect lenses. When applied to the retina, efficient coding predicts that the pervasive spatial correlations in the visual input are removed before retinal signals are sent to the brain.

In the lab of Markus Meister, I performed the first direct experimental test of this hypothesis by recording from the retinal output neurons while stimulating them with naturalistic movies. We found that the retina does remove correlations, but for the wrong reasons, namely by a nonlinear mechanism that isn't covered by the existing theory (Figure 2). This led us to generalize the efficient coding framework to adopt a more realistic description of neural signaling, revealing that the very sparse firing under natural stimuli can now be understood as optimal in the context of efficient coding.[2]

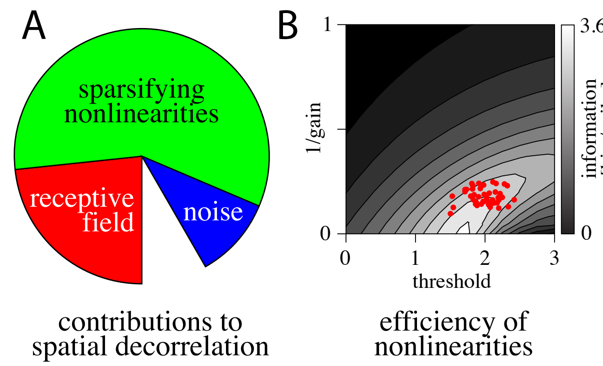

Figure 2. Retinal ganglion cells do remove correlations present in natural scenes, but (A) much more by nonlinear mechanisms than by the receptive field surround. The full circle represents the correlations present in a natural stimulus, the missing wedge represents correlations remaining after retinal processing, and the colored wedges indicate the fraction of decorrelation caused by different mechanisms. (B) Information transmission rate (contours) at a given mean spike rate depends on the type of nonlinearity. Measured nonlinearities (red dots) transmit information near-optimally.

EXACT FEATURE PROBABILITIES IN IMAGES WITH OCCLUSION. Some of the most behaviorally relevant features of natural scenes are edges caused by boundaries of objects as they occlude more distant things. Yet vision science has largely neglected to model images with occlusion. Instead, most models pretend for convenience that all objects in the world are transparent and just add together to make an image. I rectified this situation by solving a mathematical model of images generated by occlusion (Figure 3). This solution uncovered reasons for several unusual properties of edge features in natural images, unifying multiple studies of natural scene statistics. The solution allows one to calculate the joint probabilities for low-level image features as well as for high-level object properties like location, size, shape, depth, and texture, which govern the structure of the visual world and thus underly our ability to see.[3] By quantifying these relationships I was able provide new benchmarks for how biological and artificial visual systems could synthesize sensory information efficiently.

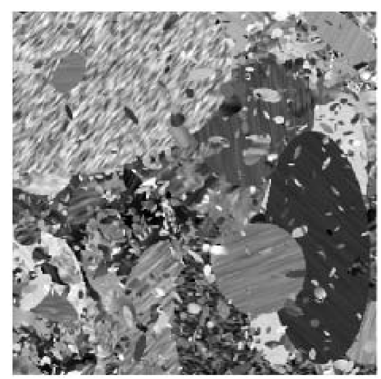

Figure 3. Example image generated by a naturalistic model with occlusion. The joint probabilities of image features and object properties like size, position, depth, and texture can be calculated exactly for this model.

LEARNING UNBELIEVABLE PROBABILITIES. Several groups have proposed that neural networks perform inference using the popular machine learning algorithm known as loopy belief propagation. This is a message-passing algorithm that disseminates local information to improve global inference. Unfortunately, the algorithm can only perform approximate inference. With Yashar Ahmadian and Ken Miller, I introduced a learning rule hoping to compensate for the approximation. However, by analyzing the learning, I proved that many desired outputs were actually 'unbelievable' — they could never be reached by belief propagation under any conditions — despite prominent claims to the contrary in the machine learning literature.[4] I then found a way to fill these gaps with a new, modified algorithm that enjoys the same advantages as belief propagation but can produce exact results even for unbelievable cases.[5] This has interesting implications for neural models based on belief propagation, since biological networks could perform dramatically better inference counterintuitively by allowing certain fluctuations in synaptic weights.

SUBOPTIMAL INFERENCE AND NEURAL VARIABILITY. Neural computations for some simple problems approach the fundamental performance limits imposed by the laws of physics, but optimal solutions to complex natural problems are largely unrealizable by neural systems. Instead, the brain must use approximations to accomplish tasks within its limited means. In a study with Jeff Beck, Wei Ji Ma, Peter Latham, and Alex Pouget, we argue that such suboptimal approximations are a dominant yet mostly unappreciated cause of behavioral variability and, under certain conditions, of neural variability.[6]

Current and future research

My goal is to develop principled models for how the brain interprets complex sensory scenes. Truly natural scenes are too complicated for this purpose, and it is often unclear how to define success quantitatively. This was my original motivation to solve the artificial, yet still naturalistic, image model described above.[3] It provides me with a powerful new tool for studying vision by isolating key properties of natural scenes and by providing a clear and calculable ground truth. This avoids some of the hardest problems of computer vision in fully natural scenes, so I can focus instead on biological mechanisms the brain uses to solve more tractable problems in complex scenes.

PERCEPTUAL INFERENCE FROM OCCLUSION: Since the occlusion image model neatly reproduces so many properties of natural scenes, one wonders whether it accurately reflects prior knowledge that humans use in visual computations. A recent study showed that artificial neural networks trained to extract reflectance from occlusion-based images misinterpret visual stimuli in ways consistent with illusory percepts.[7] I plan to address whether such illusions are interpretable as perceptual inference with this prior. Preliminary results show that the Craik-O'Brien-Cornsweet illusion can be explained this way: Reflectance optimally inferred from retinally filtered occlusion images replicates the illusion's measured dependence on contrast and size.[8] As a second example, the empirical statistics that account for how humans identify contours across occluders[9] are reproduced by the image model, implying that the model captures the prior knowledge humans use for that task.[3] By parametrically varying the image model or limiting the information available for inference, I can identify which properties are crucial for known illusions, and generate predictions for human behavior on related perceptual tasks.

NEURAL CIRCUITS FOR CONTEXTUAL INFERENCE: Researchers in machine learning have developed principled methods for statistical inference and models for which these methods are tractable. However, biological facts limit what implementations are possible in the brain. I aim to combine principles from machine learning with biological constraints to construct neurally plausible models of probabilistic inference.

To manipulate probabilistic information, the brain must first represent it. There are competing proposals about how this might happen, and there may be multiple mechanisms. One simple possibility is that the activity of a neuron is directly related to the probability of a preferred feature. Neural interactions then reflect how context influences each of those probabilities; the function of the whole neural network is to synthesize all the possible contextual effects. Unfortunately, exactly accounting for context is intractable for even modest-sized networks. Approximate algorithms have been proposed, but many of these are ill-suited to biological implementations. I showed that a variation of one algorithm, belief propagation, can be implemented biologically if each neuron has depressing synapses, a saturating output nonlinearity, and synaptic weights related to the joint statistics of the preferred features.[10] Another neurally plausible algorithm based on statistical physics[11] performs more accurately for some connectivity patterns, at the additional cost of a synaptic efficacy that varies with postsynaptic activity.[12] Ongoing efforts aim to identify biophysically realistic mechanisms that possess these properties, and test the quality of the resultant inference. Additionally I am studying how contextual inference can be implemented in another representation of probabilities, the probabilistic population code.

NEURAL CIRCUITS FOR INFERRING OBJECT BOUNDARIES: I have begun to apply these neural models to perceptual inference tasks, with promising preliminary results. A natural scene has many properties the brain must infer, so I decided to investigate the location of object boundaries. A luminance edge provides some evidence for an object boundary, but might instead indicate merely part of a texture. Remarkably, I derived that the relative probability of an object boundary is equivalent to a well-known 'energy model' of complex cells in V1.[3] This neural computation uses only feedforward sensory evidence from a small image patch. Lateral connectivity provides an opportunity to use context to improve the estimates of object boundaries. Neurons in V1 with one orientation preference are connected mostly to other neurons with similar preferences, which is hypothesized to help integrate information along contours. In the neural version of belief propagation described above, interactions between neurons are simply related to the joint probability of the neurons' preferred features, so connectivity can be predicted directly from image statistics. The result qualitatively resembles the pattern of synaptic couplings observed in the brain. Combining the connectivity and neural dynamics, simulations revealed that contextual information does improve the encoding of object boundaries (Figure 4). This success in relating neural circuitry to statistical inference invites us to apply this method to circuits deeper in the brain. For instance, in areas V2 and V4, I will use the occlusion model to predict lateral connectivity that improves the quality of inference about relative depth and shape.

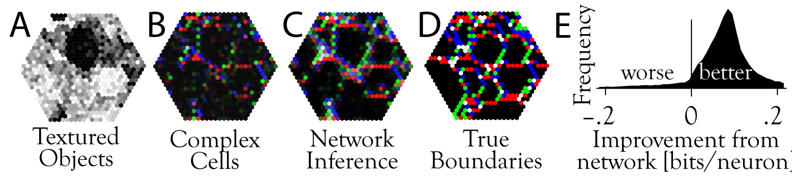

Figure 4. Model for neural inference of object boundaries. An image of random textured objects (A) gives feedforward input to model complex cells encoding the probability of an object boundary (B). Responses of model complex cells are plotted with different colors for different preferred orientations. The neurons are coupled with synaptic strengths related to the joint probability of those two object boundaries, and dynamics are given by a neurally plausible approximation to belief propagation. After a few iterations of these dynamics (C), the neurons generally have a better estimate of the true boundaries (D) than they had from feedforward evidence alone (B). The network produces an accuracy gain of about 0.1 bits per neuron (E), an improvement of approximately 10%.

FEEDBACK AND HIERARCHICAL INFERENCE: The computational roles of feedback between brain regions remains poorly understood despite its clear importance. One interesting hypothesis suggests that feedback implements inference in a hierarchical world.[13] I plan to explore the computational function of feedback in a simplified world that isolates the purely hierarchical aspects of stimuli. I will derive optimal mathematical algorithms needed for hierarchical inference, and identify key properties that neural networks need to implement them. This can be used to predict the properties of feedback connectivity, and to evaluate whether experimentally measured feedback modulates responses in a manner consistent with hierarchical inference. Ultimately, I want to integrate studies of lateral connectivity with studies of hierarchical connectivity to understand the collective function of neural circuits that span brain regions.

NEURAL REPRESENTATIONS OF PROBABILITY: Theoretical models provide biologically plausible representations of low-dimensional stimulus probabilities, but, crucially, high-dimensional ones have been more difficult to model. Recently I showed that sparse, high-dimensional probability distributions can be compressed into the average firing rates of remarkably few randomly connected threshold neurons. This representation requires exponentially fewer neurons than a complete listing of possible states, and functions without any learning using only extremely simple, neurally plausible circuitry.[14] The sparse distributions for which this representation is effective emerge naturally as sensory-conditioned probabilities generated in the hierarchically structured world described above. I plan to investigate how this random representation could facilitate probabilistic computations in a hierarchical environment, and to predict neural and synaptic properties based on these findings.

LEARNING AND INFERENCE: I assumed above that neural circuitry was already wired to exploit prior knowledge. An important issue is how this structure can be learned from the environment. Learning and perception are both inference processes, just at different timescales and with different mechanisms, so techniques used to understand them are closely related. I am interested in applying principles from machine learning to explore how synaptic plasticity allows a network to learn sensory representations and recurrent connectivities useful for inference.

OUTLOOK: Success in these projects would have far-reaching implications for neuroscience and medicine by relating single-neuron and circuit properties to perception and behavior. We could then explain how genetic or drug-induced differences in neural properties disturb perception, learning and inference, as associated with diseases like schizophrenia and autism. Ultimately this would help specify traits for pharmaceutical and behavioral therapies.

1. Pitkow X, Sompolinsky H, Meister M (2007). Visual acuity in the presence of fixational eye movements. *PLoS Biology*, **5**(12): e331.
2. Pitkow X, Meister M (2012). Decorrelation and efficient coding in retinal ganglion cells. *Nature Neuroscience*. doi: 10.1038/nn-3064
3. Pitkow X (2010). Exact feature probabilities in images with occlusion. *Journal of Vision* 10(14): 42.
4. Wainwright M, Jordan M (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* 1: 1–305.
5. Pitkow X, Ahmadian Y, Miller KD (2011). Learning unbelievable marginal probabilities. *Adv. in Neural Information Processing Systems*.
6. Beck J, Ma WJ, Pitkow X, Latham P, Pouget A (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74(1): 30–39.
7. Corney D, Lotto RB (2007). What are lightness illusions and why do we see them? *PLoS Comp Bio* 3(9): e180.
8. Pitkow X (2009). Inference of object attributes from local image features caused by occlusion. *Frontiers in Systems Neuroscience Conf Abs: Comp and Sys Neurosci*.
9. Geisler W, Perry J, Super B, Gallogly D (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41: 711–724.
10. Pitkow X, Ahmadian Y, Miller KD (2010). The value of lateral connectivity in visual cortex for interpreting natural images. *Comp and Sys Neurosci abstract*.
11. Thouless DJ, Anderson PW, Palmer RG (1977). Solution of 'solvable model of a spin glass'. *Philos. Mag.* 35: 593–601.
12. Pitkow X, Ahmadian Y, Miller KD (2011). Approximate inference by neural circuits. *Nature Precedings: Computational and Systems Neuroscience*.
13. Lee TS, Mumford D (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America* 20(7): 1434–48.
14. Pitkow X (2012). Compressive representation of high-dimensional probabilities. *Comp and Sys Neurosci abstract*.

Teaching Statement

xaq pitkow

My goal in the classroom is for students to learn concepts, models, and skills. These three form a hierarchy that frames my view of teaching.

Concepts. At the top of the hierarchy are concepts. These are the most important things we can teach students, since they are the foundation of understanding. In the life sciences the overarching concept is evolution, but another unifying concept is information, our quantitative change in uncertainty after we measure something. Every living organism from bacterium to human expends energy to acquire, manipulate, and exploit information. Sensing and learning internalize it, computation selects from it, memory stores it, motor systems use it to act. Many other important ideas in brain science — inference, decision-making, prediction, adaptation, learning — all relate to some aspect of information processing. Because biological systems can seem so haphazard, it is extremely valuable to have big-picture concepts to make sense of the complexity. They provide mental frameworks to organize facts and establish expectations about how a system will behave. This is useful even when those expectations are not met. One of the joys of science is pondering unexplained mysteries, and we can use even those examples where available concepts fall short to inspire students to ask new questions, formulate new hypotheses, and ultimately carry out new research.

Models. At the second level, models enable us to apply concepts to concrete situations. For example, natural sensory data are strongly correlated over time, so information is often found in temporal changes. Sensibly, the brain is tuned to such changes, a fact that is easily demonstrated by sensory illusions. An elementary model of neural processing that extracts the information in the changes does a good job of describing many observations, such as the transient response properties of many neurons, and demonstrates a clear application of the concept of information processing. In my classroom, students will develop a familiarity with models, and learn to evaluate their strengths and weaknesses. Students will practice evaluating models by doing small guided projects and by critiquing published models in short papers or in classroom discussions.

Skills. At the third level are the special skills are needed in the study of complex systems. We can only understand complex emergent behaviors by quantitatively accounting for the interactions of many parts. To make progress, we often need to find appropriate simplifications that eliminate many details while preserving the behaviors of interest, but even these simplified models can be quite complicated. Students thus need mathematical skills to handle them, computer skills to simulate them, estimation skills to guide decisions about which model details are crucial and which can be safely ignored, visualization skills to reveal comprehensible low-dimensional patterns in their high-dimensional behaviors, and statistical skills to judge how well they account for data.

Skills are acquired through experience, so students will develop their quantitative skills on concrete problems. Some problem sets will isolate specific techniques; others will pose conceptual questions for which students must use critical thinking skills to identify the relevant concepts and models. Solving problem sets in groups can help refine these quantitative skills and simultaneously develop communication skills needed to work in the collaborative environments in which science is often done.

Good communication skills are also necessary for publishing effective papers and delivering successful talks. Students benefit from guidance and practice in presenting concepts, models, and data clearly. Oral presentations are an important opportunity for this. Classroom poster sessions are another practical way of teaching students to be clear and focused when explaining; they also increase the amount of feedback each student gives and receives. Pedagogical studies in the quantitative sciences[1] found that peer-to-peer discussion was a highly effective way of learning concepts, and poster sessions and group problem-solving sessions are intensive exercises with precisely that emphasis.

Possible courses. I would enjoy teaching introductory and advanced courses in my chosen field, computational neuroscience. I would also like to teach classes in systems and cognitive neuroscience, which would appeal to similar concepts but focus more on experimental results and critical thinking and a bit less on mathematical models. I would be happy to teach interdisciplinary courses like complex systems or machine learning, as well as statistics or mathematical methods in the brain sciences.

Inclusivity. Interdisciplinary courses attract students with diverse backgrounds, which presents a challenge: how to bridge the substantial differences in scientific language and knowledge. I've experienced this challenge most recently in a computational neuroscience reading group I started at the University of Rochester. This has gathered together people from neuroscience, cognitive science, computer science, physics, and math. Since I have a broad background with some experience in those fields, I can facilitate discussion by translating concepts across disciplines. Even when audience members have a shared background, however, it can be difficult to ensure everyone is engaged with the material. I strongly encourage questions during lectures and talks in order to foster more active participation, promote debate, and build student confidence. I also strive to create time for soliciting questions and comments from less vocal students.

Assessment and self-assessment. Recently I was discussing science pedagogy with a senior professor renowned for his thoughtful and effective teaching. I told him I believed I was a good teacher. He asked me, "How do you know?" Of course I could point to certain concrete things, but the simple question made me appreciate that measurable objectives are important not just for evaluating students but also for evaluating myself as a teacher. To assess conceptual understanding, one interesting approach to teaching larger undergraduate classes uses carefully designed, multiple-choice concept questions as the subject of short in-class peer discussions and brief votes.[1] The discussions are a valuable interactive teaching tool, and the votes provide instant, objective feedback about student understanding, enabling a teacher to dynamically adapt to the current needs of the students. By experimenting with innovative ideas like that I can continually refine my teaching abilities and deliver deeper understanding to my students.

1. Mazur E (1997). *Peer Instruction: A User's Manual*. Prentice Hall, NJ.

References
xaq pitkow

CONTACT INFORMATION

Markus Meister — PhD thesis advisor
Professor of Molecular and Cellular Biology, Harvard University
52 Oxford Street, Northwest Building, Room 209, Cambridge MA 02138
meister@mcb.harvard.edu
617.496.8301

Ken Miller — postdoctoral advisor
Professor of Neuroscience, Physiology and Cellular Biophysics, Columbia University
co-Director, Center for Theoretical Neuroscience
1051 Riverside Drive, Unit 87, Kolb Research Annex 763, New York NY 10038
ken@neurotheory.columbia.edu
212.543.5238

Larry Abbott — postdoctoral advisor
Professor of Physiology and Cellular Biophysics, Columbia University
co-Director, Center for Theoretical Neuroscience
1051 Riverside Drive, Unit 87, Kolb Research Annex 759, New York NY 10038
lfabbott@columbia.edu
212.543.5070

Alex Pouget — postdoctoral advisor
Associate Professor of Brain and Cognitive Sciences and Bioengineering, University of Rochester
Meliora Hall, Room 402, University of Rochester, Rochester NY 14627
alex@bcs.rochester.edu
585.275.0760