

STATEMENT OF RESEARCH INTERESTS

When in graduate school I decided to focus my research on the physics of living matter, it struck me that some of the most profound and intriguing biophysical questions can be addressed in the study of learning in biological systems. One may even take the ability to learn one's environment, to respond to it actively, and to exploit it to one's benefit as a necessary component in a definition of life. Thus learning based on experience (variously known as information processing, sensing, adaptation, . . .) happens on all scales in biology, and it shows a remarkable degree of universality.

For example, at the molecular scale, the Lac operon in *E. coli* learns the lactose concentration to produce β -galactosidase (the lactose-metabolizing enzyme) in proper quantities. Similarly, in the sensory system, the retinal phototransduction cascade uses the information in the arrivals of photons to learn the instantaneous light intensity and thus the current visual scene. Additionally, it also learns the ambient light level (to adapt to it) and the temporal correlations (to estimate the self-motion through the world). On the scale of cellular (neuronal) networks, learning, memory, and adaptation are well-known. At yet larger scales, experiments on rodents are revealing how they learn and respond to changes in their environments; this is a simple, albeit quantifiable, example of the general phenomenon we call "learning" in everyday life. Finally, we may also view evolution as an example of learning, where entire species adapt to the world by means of natural selection.

Since learning is such a universal part of biology, it is natural to test the hypothesis that much of the design and function of living organisms results from optimizing learning in different contexts. This hypothesis allows us to attack many seemingly unrelated phenomena on the same footing and may lead to understanding of whether some biological features are accidental or are constrained uniquely by requirements of optimality for certain natural learning tasks. Moreover, the hypothesis is testable experimentally and, since learning can be formulated rigorously in terms of statistical field theories [1–3], it brings the powerful intuition and the mathematical apparatus of theoretical physics to a wide class of biological problems. Most of my publications over the past five years, and the main research goals I have set for the future, have revolved around testing the validity of the optimal design hypothesis and using it to sharpen the understanding of biological complexity.

Testing this hypothesis requires a clear and quantitative definition of what optimality of learning means. We have investigated this extensively in many contexts [1, 2, 4, 5]. As a result, I am now collaborating with Leslie Osborne (UCSF) to design a series of experiments that will test if the pursuit behavior in the primate visual system is optimal in learning the trajectory of the target and following it. Further, understanding which ingredients are necessary to succeed in learning, I am searching for biological implementations of them. To that end, in collaboration with Chris Wiggins (Columbia), I am analyzing (analytically and *in silico*) small genetic circuits to find better learners among them. Analysis of the whole genome genetic regulatory networks will reveal if they are built mostly from these subunits and are otherwise organized to learn optimally. This collaboration is funded by a "Quantitative Systems Biotechnology" grant (a continuation of our earlier SGER) from the NSF, of which I am the Co-PI.

In a number of ways, this study of general fundamentals and limitations of learning has led to specific advances in applications, and resulted in efficient solutions for some of the most difficult machine learning problems. One example is the information theoretic approach to feature selection—compressing data while preserving the relevant information in them—which promises to be an indispensable technique for many theoretical biology applications. Further, building on the ideas by Ma, we introduced a novel entropy estimation algorithm [6, 7], which we are using to analyze precision and variability in the neural code in the fly visual system [8, 9] and to reverse engineer transcriptional regulatory networks [10]. Other researchers are applying our method to different applications, ranging from linguistics and bioinformatics to control theory and dynamical systems.

As a third example, insights from learning theory led us to use the analogy between communicating nodes and spin glasses on random lattices to create a comprehensive theory of information processing in networks [11, 12], which has eluded the field for years. The approach has already proven itself fruitful by leading to one of the best currently available methods for reconstruction of transcriptional regulatory networks in mammalian and other cells using microarray expression data [10, 13–15]. Further, these techniques pave a way towards creation of tools for phenomemo-

logical simulation of coarse-grained description of responses of the networks to perturbations, including diseases (such as cancer), drugs, or experimental interventions. Such problems are very challenging, but also hint at substantial rewards in medical applications. Thus one of my principal research goals over the next few years is to further develop the technique: there is a need to advance beyond the standard Bethe and Kikuchi approximations and also to understand if focusing only on networks that process information and learn efficiently aids the reconstruction.

When discussing optimal information processing, my collaborators and I have noticed [1] that, since action takes time, only the part of the stimulus that contains information about what the world will be in the future should guide the organism's response. One may state quite rigorously that utility of learning is related to maximizing, extracting, and effectively representing useful, *predictive* information in the signal, which is the mutual information between its past and the future. We have analyzed this quantity theoretically [1, 5, 16], and, with Rob de Ruyter van Steveninck (Indiana), I am now testing if fly neurons, indeed, extract predictive features from a stimulus. Further, I showed that in the vertebrate photoreceptor, where standard considerations fail to explain the response time and the gain adaptation to the ambient light intensity, maximizing the predictive information is consistent with experiments [17]. A similar analysis is now being performed on small genetic circuits in collaboration with Mike Wall and Bill Hlavacek (LANL). Further, Randy Gallistel (Rutgers) is experimentally testing if maximizing predictive information by foraging rats is the origin of the paradigmatic matching laws in cognitive psychology. Additionally, we have designed these experiments to specifically answer if some ubiquitous animal behavioral patterns, such as the response hysteresis, sudden learning, critical periods, power law spectra of fluctuations in responses, can be explained by a field theoretic model of evolution of the animal's belief about the world, where these patterns emerge as manifestations of phase transitions.

Since a multitude of other learning systems are waiting to be examined, and since I believe that efficient learning and extraction of predictive information are important principles behind biological complexity, I expect this line of research to play the most prominent role in my nearest future. However, I also have interests in more traditional problems in biological and statistical physics. These include studying formation of direction sensitivity in the visual cortex, which I believe can be well-modeled in terms of the Kosterlitz-Thouless theory; identifying order parameters in critical systems by analyzing the mutual information between two large parts of the system (recall that order parameters are what carries information over large distances); analyzing cellular automata by means of the renormalization group; and others.

A critical ingredient for the success of this research program is being in a physical environment that is unusually biologically sophisticated. A physics appointment at the Indiana University associated with the Biocomplexity Institute would fit this requirement perfectly.

REFERENCES

1. W. Bialek, I. Nemenman, and N. Tishby, "Predictability, complexity, and learning," *Neur. Comp.*, vol. 13, pp. 2409–2463, 2001. arXiv: physics/0007070.
2. I. Nemenman and W. Bialek, "Occam factors and model-independent Bayesian learning of continuous distributions," *Phys. Rev. E*, vol. 65, no. 2, p. 026137, 2002. Also appeared at *NIPS'00*. arXiv: cond-mat/0009165.
3. C. Wiggins and I. Nemenman, "Process pathway inference via time series analysis," *Experim. Mech.*, vol. 43, no. 3, pp. 361–370, 2003. arXiv: physics/0206031.
4. T. Holy and I. Nemenman, "On impossibility of learning in a reparameterization covariant way," Tech. Rep. NSF-KITP-03-123, KITP, UCSB, 2002.
5. I. Nemenman, "Fluctuation-dissipation theorem and models of learning," Tech. Rep. NSF-KITP-04-20, KITP, UCSB, 2004. arXiv: q-bio/0402029. Submitted.
6. I. Nemenman, F. Shafee, and W. Bialek, "Entropy and inference, revisited," in *Advances in Neural Information Processing Systems 14* (T. G. Dietterich, S. Becker, and Z. Ghahramani, eds.), (Cambridge, MA), MIT Press, 2002. arXiv: physics/0108025.

7. I. Nemenman, "Inference of entropies of discrete random variables with unknown cardinalities," Tech. Rep. NSF-ITP-02-52, KITP, UCSB, 2002. arXiv: physics/0207009.
8. I. Nemenman, W. Bialek, and R. de Ruyter van Steveninck, "Entropy and information in neural spike trains: Progress on the sampling problem," *Phys. Rev. E*, vol. 69, p. 056111, 2004. arXiv: physics/0306063.
9. I. Nemenman, W. Bialek, and R. de Ruyter van Steveninck, "Neural coding of natural stimuli: Uncovering the information at high time resolution," In preparation.
10. A. Margolin, I. Nemenman, K. Basso, U. Klein, C. Wiggins, G. Stolovitzky, R. Dalla Favera, and A. Califano, "ARACNE: An algorithm for reconstruction of genetic networks in a mammalian cellular context." arXiv: q-bio.MN/0410037, 2004. Submitted.
11. I. Nemenman, "Information theory, multivariate dependence, and genetic network inference," Tech. Rep. NSF-KITP-04-54, KITP, UCSB, 2004. arXiv: q-bio/0406015.
12. I. Nemenman and N. Tishby, "An axiomatic approach to the theory of information processing in networks," In preparation.
13. A. Margolin, I. Nemenman, C. Wiggins, G. Stolovitzky, and A. Califano, "On the reconstruction of interaction networks with applications to transcriptional regulation." arXiv: q-bio.MN/0410036, 2004. Submitted.
14. K. Wang, N. Banerjee, A. Margolin, I. Nemenman, K. Basso, R. Dalla Favera, and A. Califano, "Conditional network analysis identifies candidate regulator genes in human B cells." arXiv: q-bio.MN/0411003, 2004. Submitted.
15. A. Margolin, N. Banerjee, I. Nemenman, and A. Califano, "Reverse engineering of yeast transcriptional network using the ARACNE algorithm," 2004. Submitted.
16. W. Bialek, I. Nemenman, and N. Tishby, "Complexity through nonextensivity," *Physica A*, vol. 302, pp. 89–99, 2001. arXiv: physics/0103076.
17. I. Nemenman, "Phototransduction cascade may be maximizing the predictive information," In preparation.

STATEMENT OF TEACHING PHILOSOPHY

I have enjoyed several opportunities to teach in my career. These include a newly created course in Computational Biology at Columbia University, short lecture series on Statistical Inference at NYU and UCSB, three summers as a faculty member for the Methods in Computational Neuroscience summer course at Woods Hole Marine Biological Laboratory, experience as a Teaching Assistant and a laboratory instructor at SFSU and Princeton, and advising undergraduate and graduate students (currently, I am advising a senior and a second year graduate student).

I genuinely like teaching, and I am open to and excited about teaching standard courses in the physics and theoretical physics curriculum (starting from basics physics and all the way to statistical mechanics and quantum field theory) based on the Department's requirements and needs. However, having worked in the areas of theoretical biophysics, complex systems, and machine learning, I will be most interested in teaching opportunities for advanced or special-topics courses in biophysics, statistical learning, and related fields. While targeted to senior undergraduate or junior graduate physics students, these courses will also be useful to mathematicians, computer scientists, biologists—all those who want to use quantitative methods to understand the design and function of biological systems. With enough interest and input from other departments (in particular, those in biomedical, chemical, and computational sciences), these courses can become an integral part of a new interdepartmental *quantitative biology* curriculum, which I will be happy to help to develop.

The following topics for courses seem to be of great overlap with my research experiences.

Introduction to Biophysical Modeling

There are many areas in biology where simple quantitative models provide clear answers to biologically important questions. These include bacterial chemotaxis, photon counting in photoreceptors, noise in biomolecular networks, protein folding, problems in adhesion, polymer and membrane dynamics, protein-DNA interactions, pattern formation, biological scalings, population dynamics, and many others. Currently, there is no single textbook covering such a wide range of modern topics, though a variety of books, reviews, and lecture notes by leading scientists in the fields are available. Thus I would like this course to be a graduate level theoretical and computational course based on original research articles. Since a crucial aspect of a graduate education is the transition from a student to an active researcher, the course will be of a seminar type, with many topics presented by the students themselves.

Statistical Physics and Statistical Inference

Many computer science and statistics departments now offer similar courses with the aim of introducing students to new data analysis methods emerging from machine learning. *The Elements of Statistical Learning* by T. Hastie et al. is a usual textbook. While hands-on experience with different algorithms is important, this is not the goal of my envisioned course. Instead I would like to make sure that the students in this undergraduate course will see the field not as seemingly unrelated computational techniques, but as a coherent subject, where similarities of the methods are emphasized, and the general requirements for any learning to be successful are studied. Thus a textbook *Information Theory, Inference, and Learning Algorithms* by D. J. C. MacKay will be used, as well as my own notes from the short courses I developed and presented at NYU and UCSB. Among others topics, we will study artificial neural networks, various coding schemes and optimization problems, and their analogs in statistical physics of disordered systems. Theoretical developments will be applied to analysis of gene expression data and to other biologically relevant problems.

Information Theory, Learning, and Cognition

This course will have some overlap with the above; however, its main goal will not be the statistical physics of learning, but rather its application to understanding how we think. I will build this course based on the lectures *Thinking about the brain* by W. Bialek (Les Houches Summer School, 2001) and on my own notes on the subject, which are emerging as a result of my collaboration with various experimental cognitive neuroscience groups. As models needed to quantify the phenomena may involve elaborate mathematics, the course will be aimed at graduate students, but it

should be accessible to advanced undergraduates as well.

Introduction to Computational Neuroscience

Computational Neuroscience is arguably the branch of biology in which quantitative methods have been the most successful in asking and answering the right questions. A course that introduces students to standard tools in neurophysiological modeling, but is mostly focused on the understanding of physical, mathematical, and experimental foundations behind such models would be of great benefit not only to students interested in neuroscience but to students interested in quantitative biology in general. Excellent textbooks such as *Theoretical Neuroscience*, P. Dayan and L. Abbott and *Spikes* by Rieke et al. are available to guide the presentation.

Introduction to Computational Biology

I will base this course on the *Computational Biology* class I am now co-teaching at Columbia. Its course description reads: "This course is designed to provide students with a broad theoretical and practical knowledge of modern computational methods as they are applied to problems in biology. Examples will be drawn from the primary literature, and topics in functional and integrative genomics will be emphasized. The course is organized around four thematic units. The first unit covers basic methods and fundamental topics, such as optimization. . . The second unit is focused on comparative genomics with an emphasis on phylogeny reconstruction. . . The third unit will cover gene expression microarray analysis. . . The final unit will cover methods used for biological network reconstruction. . . Among the recurring themes in the course will be the impact of noise and uncertainty and the impact of experimental considerations on data analysis and interpretation." This course would have to be modified to complement, but not overlap with similar courses offerer at Indiana University.

Special Topics in Statistical Physics and Applied Mathematics

In addition to biologically oriented courses, I would also be interested in (co-)teaching an advanced statistical physics/applied mathematics class. Such graduate (or advanced senior) special topics course will not be comprehensive, and it will focus on traditional mathematical and physical topics that, due to their relevance to various biological problems, fall within my sphere of expertise. In particular, I will discuss some well known problems in pattern formation (continuous systems and cellular automata, deterministic and stochastic), stochastic processes in equilibrium and nonequilibrium settings with an emphasis on biochemical reactions networks, phase transitions in optimization problems, information-theoretical analysis of systems at phase transitions, and others. Books like *Stochastic Processes in Physics and Chemistry* by Van Kampen, modern review articles, and freely available simulation tools will be used.