Statement of Research Interests

Alexander G Dimitrov

September, 2008

My main research interests involve the study of neural information processing, neural coding and information representation in biological systems. In particular I am interested in understanding information processing functions of neural ensemble activity and the biological mechanisms through which these functions are implemented. My current research concentrates on three basic aspects related to these issues: developing analytical tools and quantitative approaches to characterizing the neural representation of sensory stimuli; studying the statistical properties of natural sensory signals and their relations to biological sensory systems; and studying structure/function relations in biophysical models of neural systems. I plan to pursue these topics in the future as well. There are two principal goals which will continue to drive my research: I) to deeply understand the principles by which biological sensory systems operate; and II) to transfer these principles to engineered systems.

Animal models that I currently use in my work are: the cricket cercal system (in collaboration with Drs. Miller and Jacobs at MSU), the visual systems of cat and macaque (in collaboration with Dr. Gray at MSU and Dr. Yen at NUS, Singapore) and rat vibrissal system (in collaboration with Dr. Azouz at Ben Gurion University and Dr. Jackie Schiller, Technion, Israel).

My original training is in quantitative disciplines (Physics, Applied Mathematics, Mathematical Neurobiology). One of the most important goals which I achieved during my postdoctoral work at Dr. John Miller's lab was to obtain sound training in neurophysiological and behavioral experimental methods and procedures. Since then I have continued to work in close collaboration with experimentalists, being closely involved in experimental design, data collection and analysis.

The mathematical tools that I apply to achieve these goals come mostly from branches of applied probability (information theory, signal processing theory, stochastic differential equations), dynamical systems theory and group theory. Many other branches of mathematics and statistics – optimization, operations research, and differential geometry being the ones I used most recently – are applicable in particular stages throughout my research

Current Research

The main direction of research I have pursued is on information processing in biological sensory systems. It involves understanding of the nature of both sensory stimuli and associated neural activity. The research is naturally partitioned in two related but independent parts - characterizing the neural code of a sensory system, and characterizing the statistical structure of sensory stimuli, relevant for the particular system.

Neural coding. One of the steps toward understanding the neural basis of an animal's behavior is characterizing the code with which its nervous system represents information:

the correspondence between environmental or internal stimuli and activity of a set of neurons. All computations underlying an animal's behavioral decisions are carried out within the context of this neural code. A determination of the neural coding schemes is an extremely important goal, due not only to our interest in the nature of the code itself, but also to the constraints that this knowledge places on the development of theories for the biophysical mechanisms underlying neural computation. Knowledge of the neural coding scheme is also an essential ingredient in the emerging field of neural prosthetics, enabling the restoration of neural function to occur without the need for the brain to adapt to the new coding schemes introduced by neural prosthetic devices.

Some central questions to be addressed when discussing neural coding are: What stimulus features are encoded in neural activity patterns? What aspects of the neural activity patterns encode that information? What are the algorithms through which information is encoded and decoded from ensemble activity patterns? What processing tasks are performed on sensory information? Considerable progress has been made by approaching these questions independently. However, independent treatment of these interconnected questions often introduces multiple assumptions that prevent their complete solution.

To address these issues, my collaborators and I recently presented an analytical approach [19, 20, 23, 26] that enables the simultaneous solution to these interconnected questions. The basis for this approach is to conceptualize a neural coding scheme as a collection of stimulus-response classes, where each class consists of a set of stimuli and a synonymous set of neural responses. The stimulus-response classes form a structure akin to a dictionary or 'codebook', with each class being represented by a set of activity patterns and their associated stimulus features. Our analytical approach enables the derivation of this neural codebook, which in turn allows any sequence of spike patterns in a neural response to be related to the corresponding set of stimulus features that elicited those responses.

The impetus behind this line of research was the realization that the functional issues that confront the early stages of any biological sensory system are similar to the issues encountered by communication engineers in their work of transmitting messages across noisy media, which is the domain of Shannon's Information Theory. With this in mind we represented the input/output relationship present in a biological sensory system as a *communication system* [19, 20, 23, 26]. This approach has been suggested before, however, all the properties that information theory assigns to this object have not been completely appreciated in the neural research literature. Even though this coding model is probabilistic, results from information theory imply that an almost deterministic relation emerges naturally on the level of clusters of stimulus/response pairs [20].

This analysis uses tools from two branches of applied probability: information theory and quantization theory. Briefly, the subject information theory is transmission and recovery of information over noisy media. Quantization theory deals with optimal reduced representations of random variables, which satisfy a fidelity criterion. The new analytic approach quantizes (clusters) the neural responses to a small reproduction set and optimizes the quantization to minimize an information-based fidelity criterion. In this way, fixing the size of the reproduction set produces the most informative approximation of the coding scheme at this resolution. This approximation can be automatically refined as more data becomes available.

To use whatever limited data set is available most efficiently, I quantize neural responses to a reproduction set of small finite size [20]. The choice of an optimal quantization is achieved through the use of the recently designed information distortion function [18, 20, 21]. This method allows the study of coarse but highly informative models of a coding scheme which can be automatically refined when more data becomes available [22, 23].

Ultimately, a simple description of the stimulus/response relation can be recovered.

The structure of sensory stimuli. The structure of the early sensory pathways in most organisms is determined in part by the structure of stimuli they perceive from the surrounding environment. It is likely that evolutionary pressures had shaped the development of their sensory systems so that they optimally process the incoming sensory data stream according to certain information-theoretic measures. I work under the assumption that the architecture in many parts of an organism's brain is determined by external signals, partially genetically and partially through neural plasticity [7, 14, 16].

My work in this area was concentrated on the visual system and area V1 of the primate visual cortex. I investigated the appearance of orientation-selective cells in cortical area V1 and the lateral connectivity between them. I studied the statistical structure of visual signals from an organism's natural environment and estimated some of their information-theoretic measures (autocorrelations, entropy, mutual information) [14, 15, 16]. Natural scenes are highly redundant according to these measures. An optimal information processing system will transform the signals in a way to optimize the representation by minimizing redundancy. Assuming that a part of the cortex performs such an optimization, I derived the structure of units and connectivity between them to achieve such a task [7, 14, 15, 16].

My current work transfers these techniques to a different sensory environment - the set of air currents which stimulate the cricket cercal sensory system. I am in the process of recording and characterizing the crickets' natural airflow environment: ambient airflows and signatures of natural predators, and relating these to the observed functional properties of the system. However, a more interesting possibility emerged during this research - that the important stimulus structures are not only the ones introduced by the environment, as the prevalent paradigm of natural stimulus studies implies, but also the ones relevant to the particular animals' lifestyle [3]. In particular, while current research on natural stimuli usually concentrates on exact representation of the incoming signal, I am one of the few researchers that is investigating the hypothesis that the incoming signal is heavily compressed (quantized), with information losses controlled by the animals' intrinsic needs. Characterizing this lossy compression distortion function [9] has been the impetus behind developing tools for dejittering [3, 13, 17, 24] and for adaptive and unequal probability sampling in biological sensory systems [1, 11].

Biophysical models of neural systems It becomes more and more evident that cellular dynamics is very important for the functionality of biological brains – the environment in which they operate has its own dynamics in many different time scales. Some of our previous results suggest that many of the cortical subsystems are operating near a bifurcation point, which enhances their response properties [27, 28]. Recent work [5, 6] also suggests that nonlinearities in the propagation of spikes may be used to optimize coding properties in the sensory system. I am currently investigating nonlinear dynamics, bifurcation structure and information processing properties of dynamical systems models of neural systems [4]. The approach I am taking combines classical deterministic systems with probability, to model both the dynamics and uncertainty of neural operations. I am using classical results that convert deterministic systems to their corresponding Fokker-Planck equations, and recent work by Paninski that provides a principled approach to finding parameters of these equations. My students and I are currently adapting these equations to the cricket cercal system and extending them to include more of the detailed knowledge available in this sensory system [4, 29, 30] .

Directions for future research

Neural coding. The above methods of uncovering neural coding schemes provide a unique opportunity for quantitative developments in our understanding of sensory systems. The quantization procedure provides an explicit probabilistic description of the neural stimulus/response relationship. We can use it to modify the initially presented stimulus to be more relevant to the functioning of the investigated sensory system. This procedure of stimulus optimization allows us to iteratively refine our knowledge of the sensory system by automatically presenting to the system results of the information-theoretic analysis. I am already heading a project in this direction [1]. My students and I are currently building a data acquisition and analysis system that can perform this analysis and steer the iterations of the stimulus in real time. In its core are statistical techniques from adaptive sampling theory [12].

The structure of sensory stimuli. I plan to continue studying the statistical structure of sensory signals from an organism's natural environment, extending beyond the visual domain into auditory, olfactory and mechanosensory modalities. Of course, optimal information processing is only one of the constraints on the performance of any brain. To obtain realistic models of neural function and activity, we have to take into account other important constraints – metabolic efficiency, wiring cost and unit failure [10]. It is conceivable that different organisms have achieved distinct adaptations to the same sensory environment by adopting different tradeoffs between these conflicting demands. I plan to study the relative importance of these constraints, and incorporate them in a set of models which can explain differences in the neuronal structure of various organisms.

The quantization formalism discussed previously allows for another set of interesting questions to be asked about biological sensory systems. Specifically, one can inquire about what details of the sensory stimulus are retained, and what are discarded at specific stages of sensory processing [9]. This is a natural question in the framework of Rate Distortion Theory, whence quantization comes from, and which can provide a consistent set of tools with which to address this questions. Some preliminary results in this direction prepared the formalism and groundwork for addressing such questions [17], and successfully tested the initial predictions in the cricket cercal sensory system [2, 3, 25]. Currently we are expanding the formalism to cover more cases of interest, and applying it to clarifying functional properties in the cat primary visual cortex [12, 24].

Biophysical models of neural systems. The functional properties that our analysis uncovers are implemented through a particular neural architecture. A detailed knowledge of the interplay between structure and function in a nervous system will increase our understanding of how the system operates as a whole. The cricket cercal sensory system is well suited for approaching these questions, since it balances simplicity and complexity to provide sufficiently general, yet manageable animal model. I plan to pursue further research in how cell dynamics and anatomical structure produce observed functional response. Currently my interests lie in modifying the probabilistic dynamics model developed by Paninski et. al into a cercal afferent model, and extending the Fokker-Planck formalism to more biophysically precise models [29, 30].

Structural and functional robustness in neural systems. Our improved understanding of neural coding has brought forth an interesting relation to another aspect of biological neural systems – robustness. One limiting case is the neural architecture of large, long-lived systems (many mammals, reptiles), where it is almost certain that some neural units will fail during the expected lifetime of the system. This affects the structural and functional organization of such systems: They have redundant architecture, where no single unit carries out much of the processing. The units are unreliable (low release probability of neurotransmitters) and most likely inexpensive to replace. The coding is most likely probabilistic. The system can recover gracefully from unit loss, as the remaining units reorganize to take over some of the tasks of the dysfunctional units. Some recent research also finds evidence for the production of replacement units.

This can be contrasted to another limiting case – of relatively small, short living systems (e.g., most insects), where the expected lifetime of neurons is likely longer than that of the system (mostly due to predation). This respectively affects the structure of their nervous system. Insects usually have fewer, identifiable neurons, with little structural redundancy. Each unit can be shown to perform substantial processing. The units and interactions between them are in general more reliable than those of mammals. On the other hand a failure of a few sensory interneurons there almost certainly leads to the demise of the system. However, redundancy is also present, albeit predominantly on a population level: for example, insects have multiple offspring which are almost genetically identical.

Preliminary results of this research were presented at a workshop on the general topic of optimization in brain design [8, 10]. I plan to pursue this research topic further. To me it has profound ramifications in understanding the functioning of neural systems. I also believe that design principles uncovered by this research can also be applied toward developing novel computing architectures.

Bio-inspired Engineering. One of my longstanding interests has been to apply my knowledge on neuronal networks in biological systems to artificial intelligence. Most of the models that are built to explain the structure and functionality of the brain are directly transferable as functional models in an intelligent agent. Biological systems had evolved over billions of years to adapt exquisitely to their environments. More often than not, modern engineering inventions have biological analogues of much higher quality and performance. Some classic examples are the sonar (evolved in bats, dolphins), reactive jet engine (squid), vision (in most vertebrates; still unmatched by machine vision).

One can safely assume that transferring knowledge of neurobiological function to engineering problems will lead to a more rapid pace of development and yield superior devices. However, most bio-inspired design today relies on an approach inverse to the normal engineering design process: a biological structure-function relation is understood and then a search is conducted for an engineering problem that may be solved with that knowledge. The usual forward approach (and the one most useful to technological advance) - an engineering problem in search of a biological solution paradigm - is today virtually unavailable. What I and colleagues have proposed recently is to remedy this deficiency by facilitating the *systematic technology transfer between biology and engineering* (grant proposal to DARPA). In the proposed new process, an engineer would define as much of the technology problem specifications and constraints as possible and then query a database of biological structure-function relations, searching for biological cases with similar specifications and constraints.

Conclusions

The study of biological neural systems poses many challenging problems, the analysis of which involves elaborate analytical and experimental techniques. As in other problems where experiments meet theory, the analysis of a complex physical system strains the current state of the art of available tools. I expect this general line of research to act as a catalyst for exciting new development in our understanding of how biological sensory processing and perception.

References

[1] Z. Aldworth, A. G. Dimitrov, and J. P. Miller. Iterative stimulus refinement uncovers intrinsic stimulus features represented by the cricket cercal sensory system. *in preparation*, 2008.

[2] Z. Aldworth, A. G. Dimitrov, and J. P. Miller. Nonlinear coding with spike patters in interneurons of the cricket cercal sensory system. *in preparation*, 2008.

[3] Z. N. Aldworth, J. P. Miller, T. Gedeon, G. I. Cummins, and A. G. Dimitrov. Dejittered spike-conditioned stimulus waveforms yield improved estimates of neuronal feature sensitivity. *J. Neurosci.*, 25(22):5323–5332, 2005.

[4] B. Bartle, G. I. Cummins, and A. G. Dimitrov. Mathematical modeling of afferent neurons in the cricket. MSU REU poster session, August 2005.

[5] J. A. Bender, A. G. Dimitrov, and J. P. Miller. Biophysical constraints on the precision of neural codes. In *Society for Neuroscience Annual Meeting*, 2001. (abstract).

[6] G. I. Cummins, S. Crook, T. Ganje, A. G. Dimitrov, and G. Jacobs. Structural and biophysical mechanisms underlying dynamic sensitivity of primary sensory interneurons in the cricket cercal sensory system. *Neurocomputing*, 52-54:45–52, 2003.

[7] A. G. Dimitrov. *Aspects of Cortical Information Processing.* PhD thesis, The University of Chicago, 1998.

[8] A. G. Dimitrov. Unit failure and circuit complexity. Workshop "Optimization and constraints in the evolution of brain design", CSHL, July 2003.

[9] A. G. Dimitrov. Lossy compression in neural sensory systems: At what cost (function)?. Methods of Information Theory in Computational Neuroscience CNS*06 Workshop, July 2006. (invited talk).

[10] A. G. Dimitrov. Why do long-living animals need larger brains? seminar, the Gatsby Computational Neuroscience Unit, UCL, London, July 2006. (invited talk).

[11] A. G. Dimitrov. What you show is what you get: sampling biases in determining biological sensory function. In *The Soceity for Neuroscience Annual Meeting*, San Diego, November 2007. (poster).

[12] A. G. Dimitrov. What you show is what you get: sampling biases in determining biological sensory function. poster, 4th International Workshop *Statistical Analysis of Neuronal Data*, Pittsburgh, PA, May 2008.

[13] A. G. Dimitrov, R. Azouz, and L. Israeli. Effects of stimulus transformations on estimated functional properties of mechanosensory neurons. *Neurocomputing*, 70:1772–1776, 2007.

[14] A. G. Dimitrov and J. D. Cowan. Spatial decorrelation in orientation tuned cortical cells. In M. Mozer, M. Joran, and T. Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9. The MIT Press, 1997.

[15] A. G. Dimitrov and J. D. Cowan. Edge detectors and texture detectors differ in their lateral connectivity. In J. Bower, editor, *Computational Neuroscience: Trends in Research 1998.* Plenum Press, 1998.

[16] A. G. Dimitrov and J. D. Cowan. Spatial decorrelation in orientation selective cortical cells. *Neural Computation*, 10(7):1779–1796, 1998.

[17] A. G. Dimitrov and T. Gedeon. Effects of stimulus transformations on the perceived function of sensory neurons. *JCNS*, 20:265–283, 2006.

[18] A. G. Dimitrov and J. P. Miller. Analyzing sensory systems with the information distortion function. In R. B. Altman, editor, *Pacific Symposium on Biocomputing 2001*. World Scientific Publushing Co., 2000.

[19] A. G. Dimitrov and J. P. Miller. Natural time scales for neural encoding. *Neurocomputing*, 32-33:1027–1034, 2000.

[20] A. G. Dimitrov and J. P. Miller. Neural coding and decoding: communication channels and quantization. *Network: Computation in Neural Systems*, 12:441–472, 2001.

[21] A. G. Dimitrov, J. P. Miller, Z. Aldworth, and T. Gedeon. Non-uniform quantization of neural spike sequences through an information distortion measure. *Neurocomputing*, 38-40:175–181, 2001.

[22] A. G. Dimitrov, J. P. Miller, Z. Aldworth, and A. Parker. Spike pattern-based coding schemes in the cricket cercal sensory system. *Neurocomputing*, 44-46:373–379, 2002.

[23] A. G. Dimitrov, J. P. Miller, T. Gedeon, Z. Aldworth, and A. E. Parker. Analysis of neural coding through quantization with an information-based distortion measure. *Network: Computation in Neural Systems*, 14:151–176, 2003.

[24] A. G. Dimitrov, M. A. Sheiko, J. Baker, and S. Yen. Effects of stimulus transofrmations on perceived functional properties of visual sensory neurons. *in preparation*, 2008.

[25] T. Ganje, A. G. Dimitrov, and J. P. Miller. Nonlinear coding with spike patters in sensory afferents of the cricket cercal sensory system. *in preparation*, 2008.

[26] T. Gedeon, A. E. Parker, and A. G. Dimitrov. Information distortion and neural coding. *Can. Math. Q.*, 10:33–70, 2003.

[27] T. Mundel, A. G. Dimitrov, and J. D. Cowan. A simple model for cortical orientation selectivity. In J. Bower, editor, *Computational Neuroscience: Trends in Research 1997*. Plenum Press, 1997.

[28] T. Mundel, A. G. Dimitrov, and J. D. Cowan. Visual cortex circuitry and orientation tuning. In M. Mozer, M. Joran, and T. Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9. The MIT Press, 1997.

[29] J. Verrette and A. G. Dimitrov. Interneuronal extensions to the L-NLIF neural coding model. MSU Student Research Celebration, April 2008.

[30] J. Verrette, A. G. Dimitrov, and J. P. Miller. Pairwise correlations in cricket cercal interneurons are significant for decoding. In *BMC Neuroscience*, volume 8, page 159. BioMed Central Ltd., 2007.