ביה"ס להנדסה ולמדעי המחשב ע"ש רחל וסלים בנין
The Rachel and Selim Benin School of Computer Science and Engineering
האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

ELSC
The Edmond & Lily Safra Center for Brain Sciences

הקרן הלאומית למדע
ISRAEL SCIENCE FOUNDATION

Information, Control, and Learning – The Ingredients of Intelligent Behavior

September 26-28, 2016

The Brindell and Milton Gottlieb Auditorium, The Rothberg Family buildings,
Rachel and Selim Benin School of Computer Science and Engineering
Edmond J. Safra Campus, Jerusalem

Talk Abstracts

Nihat Ay – Max Planck Institute for Mathematics in the Science

Mathematical Aspects of Embodied Intelligence

I will present recent results on the design of embodied systems with concise control architectures, formalizing the notion of "cheap design" within the field of embodied intelligence. This notion highlights the fact that high behavioral complexity, seen from the external observer perspective, does not necessarily imply high control complexity. This complexity gap is a result of two different frames of reference, which is closely related to Uexküll's Umwelt concept. If time allows, I will present a measure-theoretic formalization of this concept and discuss its implications .

References:

1. N. Ay. Geometric Design Principles for Brains of Embodied Agents. KI - Künstliche Intelligenz (2015) 29: 389-399. doi: 10.1007/s13218-015-0382-z .
2. N. Ay, W. Löhr. The Umwelt of an Embodied Agent - A Measure-Theoretic Definition. Theory in Biosciences (2015) 134: 105-116. doi: 10.1007/s12064-015-0217-3 .
3. G. Montúfar, K. Ghazi-Zahedi, N. Ay. A Theory of Cheap Control in Embodied Systems. PLoS Computational Biology (2015) 11(9): e1004427. doi: 10.1371/journal.pcbi.1004427.
4. G. Montúfar, N. Ay, K. Ghazi-Zahedi. Geometry and Expressive Power of Conditional Restricted Boltzmann Machines. Journal of Machine Learning 16 (2015) 2405-2436 .

Link to the project site at the Santa Fe Institute: http://www.santafe.edu/research/embodied-intelligence/

Naama Brenner, Technion – Israel Institute of Technology

Exploratory adaptation in random networks

Biological systems exhibit well-defined and repeatable responses, but can also explore, improvise and generate new functionality in the face of unforeseen conditions. At the level of populations, exploration is implemented out by variation and selection. At the level of the single cell or organism, exploratory behavior has been observed experimentally, but a theory or organizing principle is lacking.

Here I will review experimental evidence for exploratory behavior that stabilizes adaptive phenotypes within the lifetime of an organism. I will then present a model of gene regulatory networks which demonstrates the capacity for such behavior. The model addresses the problem of convergence of temporal exploration in high-dimensional spaces, and shows that such convergence is non-universal and depends on network properties. The results suggest a deep connection between cellular adaptation and learning in neural systems.

Andrea Cavagna, Institute for Complex Systems, Rome

Information transfer and behavioral inertia in starling flocks

Collective decision-making in biological systems requires all individuals in the group to go through a behavioral change of state. During this transition fast and robust transfer of information is essential to prevent cohesion loss. The mechanism by which natural groups achieve such robustness, however, is not clear. I will present an experimental study of starling flocks performing collective turns and show that information about direction changes propagates across the flock with a linear dispersion law and negligible attenuation, hence minimizing group decoherence. These results contrast starkly with present models of collective motion, which predict diffusive transport of information. Building on spontaneous symmetry breaking and conservation-law arguments, I will formulate a theory that correctly reproduces linear and undamped propagation. Essential to this framework is the inclusion of the birds' behavioral inertia. The theory not only explains the data, but also predicts that information transfer must be faster the stronger the group's orientational order, a prediction accurately verified by the data. These results suggest that swift decision-making may be the adaptive drive for the strong behavioral polarization observed in many living groups.

Kenji Doya, Okinawa Institute of Science and Technology Graduate University

Neural Circuit Mechanisms of Mental Simulation

The basic process of decision making can be captured by learning of action values according to the theory of reinforcement learning. In our daily life, however, we rarely rely on pure trial-and-error and utilize any prior knowledge about the world to imagine what situation will happen before taking an action. How such "mental simulation" is realized in the circuit of the brain is an exciting new topic of neuroscience. Here I report our recent works with functional MRI in humans and two-photon imaging in mice to clarify how action-dependent state transition models are learned and utilized in the brain.

Amir Globerson, Tel Aviv University

Variational Conditional Probabilities

Predicting the label Y of an object X is a core task in machine learning. From a probabilistic perspective, this involves reasoning about conditional probabilities p(y|x). However, it is hard to obtain reliable estimates for these probabilities. Here we show how to obtain lower and upper bounds on p(y|x) given statistical information, and show how it can be used within various learning setups. We also extend this formulation to the structured case, where y can be multivariate.

Bert Kappen, Radboud University and University College London

Integrating control, inference and learning. Is it what the brain does?

Intelligent systems, whether natural or artificial, must act in a world that is highly unpredictable. To plan actions with uncertainty is a stochastic optimal control problem. However, there are two fundamental problems: the optimal control solution is intractable to compute and intractable to represent due the non-trivial state dependence of the optimal control. This has prevented application of stochastic optimal control theory to robotics or as a model for the brain so far.

The path integral control theory describes a class of control problems whose solution can be computed as an inference computation through Monte Carlo sampling. The sampling can be made more efficient by adaptive importance sampling. This defines a recursive learning problem, where a better importance sampler is learned from self-generated data. I formalize the intuitive notion that the efficiency of the importance sampling is related to the proximity of the sampling control to the optimal control. Secondly, I show how parametrized feed-back control functions can be estimated using the cross entropy method . I finally discuss how these ideas can be used as an abstract model for sensorimotor control.

Máté Lengyel, University of Cambridge

Neural variability and sampling-based probabilistic representations in the visual cortex

Neural responses in the visual cortex are variable, and there is now an abundance of data characterizing how the magnitude and structure of this variability depends on stimulus attributes. Current theories of cortical computation fail to account for these data: they either ignore variability altogether, or only model its unstructured Poisson-like aspects. We develop a theory in which the cortex performs probabilistic inference such that population activity patterns represent statistical samples from the inferred probability distribution. Our main prediction is that perceptual uncertainty is directly encoded by the variability, rather than the average, of cortical responses. Through direct comparisons to previously published data as well as original data analyses, we show that a sampling-based probabilistic representation accounts for the structure of noise, signal, and spontaneous response variability and correlations in the primary visual cortex. These results suggest a novel role for neural variability in cortical dynamics and computations.

Shie Mannor, Technion – Israel Institute of Technology

Regularization and Robustness in Reinforcement Learning

In many sequential decision problems all that we have is a record of historical trajectories. Building dynamic models from these trajectories and ultimately sequential decision policies may result in much uncertainty and bias. In this talk we consider the question of how to create control policies from existing historical data and how to better sample trajectories so that future control policies would be better. This question has been central in reinforcement learning in the last decade if not more, and involves methods from statistics, optimization, and control theory.

We will discuss two possible remedies to uncertainty in sequential decision problems: building robust policies and using risk measures such as the conditional value-at-risk as the objective to be optimized rather than the ubiquitous expected reward. We consider the complexity and efficiency of evaluating and optimizing risk measures and solving robust Markov decision processes. Our main theme is that considering risk is essential to obtain resilience to model uncertainty and model mismatch. Alternatively, one can obtain this resistance by solving robust Markov decision processes directly. We discuss some of the issues pertaining to solving robust Markov decision processes.

Ron Meir, Technion – Israel Institute of Technology

Optimal Neural Codes for Estimation and Control

Optimal neural codes have been widely studied in the field of neuroscience in the context of perception, suggesting explanations of encoding in terms of optimal signal representation. However, in general settings perception subserves action selection leading to behavior through the perception-action loop. Motivated by recent results from motor control, we argue that encoding strategies should depend on the task for which they are being used. Given the different requirements of purely perceptual (e.g., object classification) and action-oriented (e.g., object manipulation) tasks, it is expected that sensory adaptation should differ within these domains. We show, based on principles of optimal estimation and control, that sensory adaptation for control differs from sensory adaptation for perception, even for simple setups. This implies, consistently with experiments, that when studying sensory adaptation, it is essential to account for the underlying task. In a broader setting, top-down effects, task-dependent and otherwise, have been shown experimentally to contribute to early sensory processing, thereby inspiring the development of a theoretical framework for such phenomena.

Sanjoy Mitter, Massachusetts Institute of Technology

Information and Entropy Flow in Estimation and Control

There are two great theories in the Systems, Communications and Control field: Information Theory and Stochastic Control. In this talk, I discuss the role of Information in Estimation Theory and Stochastic Control. There are ``dynamical'' analogues of source coding and rate distortion theory, as well as channel coding in stochastic control theory, where the concepts of directed information and transfer entropy play a central role. I give an information-theoretic view of Kalman filtering and its nonlinear generalizations. I illustrate the role of directed information and transfer entropy in stochastic control by considering the problem of extracting the maximal amount of work from a noisy, electrical circuit acted upon by a Maxwell Demon.

Joint work with Nigel Newton (University of Essex, UK), Henrik Sandberg (KTH, Sweden), Jean-Charles Delvenne (UCLouvaine, Belgium) and Takashi Tanaka (KTH, Sweden)

Israel Nelken, The Hebrew University of Jerusalem

Information Theoretic Bounds and Brains: Prediction and Planning

We would like to use the general concepts surrounding perception-action loops in order to understand brain responses and behavior of real animals in (almost) natural environments. I will discuss here two simple test cases, one at the sensory side of the loop, the other in the control/action side of the loop, in which the application of information bounds provides explicit solutions that can be usefully compared with experimental results. The first case involves stimulus-specific adaptation (SSA), the reduction in the responses to a repeatedly-presented sound ('standard') that does not fully generalize to another, rare, sound ('deviant'). We interpret SSA as the expression of a prediction error signal. Using the Information Bottleneck principle, we show that the neuronal responses are consistent with a prediction error signal that is based on an internal representation that has a long memory (>10 stimuli back) which is coarse but still almost optimally predictive. Thus, information bounds allow us to estimate interesting properties of the neuronal code. The second case involves the application of information-limited control ideas to learning in the Morris Water Maze (MWM). We developed a simple LQG control model for the mouse in the MWM, and solved it explicitly with information constraints on the swimming paths. The model provides two quantifiers for actual swimming paths: one is 'value', which is closely related to the swimming time to platform, the standard quantifier of behavior in the MWM. The other is 'complexity', which quantifies the deviation of the swimming paths from the no-control distribution. While the value continuously increases during training, the complexity may be a non-monotonic function of training session, is sensitive to genetic background and mouse gender, and may be a new, sensitive quantifier of behavior in the MWM. In the future, we will use a full model of the perception-action loop in order to study animal behavior in complex, but well-controlled, environments.

Daniel Polani, University of Hertfordshire

Informational Drivers of Cognition

In the last years, (Shannon) information increasingly established itself as an essential concept for understanding the processes of decision-making and cognition. It has now grown out of being considered merely a "glorified correlation measure" and has become a language in which cognitive problems can be expressed and understood. In addition, it provides tools for the prediction of phenomena and is sufficiently operational to permit construction of cognitive models in a principled way. Furthermore, indications that biological organisms might prefer informationally optimal behaviors provide opportunities to make formally founded conjectures about possible adaptation pathways in biology. The talk will present representative examples of using information theory in these contexts, amongst others, empowerment for the generation of intrinsic motivation drives.

Alexandre Pouget, University of Geneva

Optimal policy for value-based decision making

Every-day decisions frequently require choosing among multiple alternatives each associated with its own reward. Previous physiological and behavioral experiments have revealed puzzling properties of human/animal decisions such as interactions among these options and time-dependent decision thresholds. Why the nervous system ought to have such properties and how they functionally relate to each other remains poorly understood. To address these issues, we derive the normative strategies for general value-based decisions with multiple options, and identify the optimal stopping-rules for value-based evidence accumulation. The resulting complex nonlinear and time-dependent decision-boundaries in a high-dimensional belief space appear intractable by nervous systems in situ. However, using geometric symmetries in those boundaries allows us to implement the optimal strategy by a remarkably simple neural mechanism, with time-dependent activity-normalization in a recurrent circuit controlled by an urgency signal. The model reveals why the nervous system requires such activity normalization and urgency signal: they allow the nervous system to implement optimal decisions under multi-alternative choices.

Yifat Prut, The Hebrew University of Jerusalem

Inhibitory control of motor timing

Proper performance of voluntary movements requires the integration of both spatial and temporal information about the ensuing movements. The timing of actions is often considered to be dictated by cerebellar output that is relayed to the motor cortex via the motor thalamus. We investigated the

mechanisms by which the Cerebellar-Thalamo-Cortical (CTC) system controls temporal properties of motor cortical activity.
We found that in primates the CTC pathway efficiently recruits motor cortical neurons in primary motor and premotor areas. Cortical responses to CTC activation were dominated by prolonged inhibition mediated by a feedforward mechanism. We further found that cortical cells that integrated CTC input fired transiently and synchronously at movement onset, when the timing of action is dictated. Moreover, when preventing the flow of information in the pathway the phasic firing at movement onset was reduced, but the preferred direction of the cells remained unchanged. These changes in neural firing were correlated with altered motor behavior: the monkeys were able to perform the task but with increased reaction and movement times.
These results suggest that the CTC system affects cortical firing by changing the excitation-inhibition balance at movement onset in an extensive network of TC-activated motor cortical neurons. In this manner, the temporal pattern of neural firing is shaped, and firing across groups of neurons is synchronized to generate transiently enhanced firing.

Reza Shadmehr, Johns Hopkins University

Encoding of action by the Purkinje cells of the cerebellum

In this talk I will present the hypothesis that representation of error anatomically organizes the Purkinje cells of the cerebellum into micro-clusters. Using single cell recordings during saccadic eye movements, I suggest that this anatomical clustering allows one to use the complex spike properties of Purkinje cells to understand how the simple spikes encode movements. The results suggest that the cerebellar cortex predicts motion of the eye via a gain-field encoding of direction and speed. I will then use this framework to examine error-dependent learning and the neural basis of the multiple timescales of motor memory.

Shai Shalev-Shwartz, The Hebrew University of Jerusalem & Mobileye

Deep Reinforcement Learning for Driving Policy

Technology for autonomous driving requires both "sensing" (understanding the environment) and "acting" (moving the car appropriately). The talk will focus on the "acting" part, which we call a "driving policy". Manually defining a driving policy is tricky mainly because it requires to negotiate the right of way with other drivers by balancing between unexpected behavior of other drivers while not being too defensive. Reinforcement learning seems like the appropriate tool for learning a driving policy. We discuss specific challenges to this domain, including learning with safety guarantees and reinforcement learning beyond MDPs.

Tatyana Sharpee, Salk Institute for Biological Studies

How invariant feature selectivity is achieved in cortex

Natural stimuli hold the key to understanding how advanced signal processing – those that make it possible for us to recognize specific people and events – occur in the brain. At present there is only cursory understanding of how detailed signal representations provided by the sensory periphery are transformed in the brain to give rise invariant forms of selectivity observed deep inside the brain, in both the visual and auditory system. Towards this goal, I will describe a set of statistical methods that can be used in conjunction with natural stimuli to probe and characterize feature selectivity and invariance of neurons deep inside the sensory circuits. I will present new insights that these methods have produces about the responses of mid-level visual and auditory neurons.

Satinder Singh Baveja, University of Michigan

Rethinking State Action and Reward in Reinforcement Learning

Over the last decade and more, there has been rapid theoretical and empirical progress in reinforcement learning (RL) using the well-established formalisms of Markov decision processes (MDPs) and partially observable MDPs or POMDPs. At the core of these formalisms are particular formulations of the elemental notions of state, action, and reward that have served the field of RL so well. In this talk, I will describe recent progress in rethinking these basic elements to take the field beyond (PO)MDPs. In particular, I will briefly describe older work on flexible notions of actions called options, briefly describe some recent work on intrinsic rather than extrinsic rewards, and then spend the bulk of my time on recent work on predictive representations of state. I will conclude by arguing that taken together these advances point the way for RL to address the many challenges of building an artificial intelligence.

Stefano Soatto, University of California, Los Angeles

Perception, Action and the Information Knot that Ties Them

I will describe a notion of Information for the purpose of decision and control tasks, as opposed to data transmission and storage tasks implicit in Communication Theory. It is rooted in ideas of J. J. Gibson, and is specific to classes of tasks and nuisance factors affecting the data formation process. When such nuisances involve scaling and occlusion phenomena, as in most imaging modalities, the ``Information Gap'' between the maximal invariants and the minimal sufficient statistics can only be closed by exercising control on the sensing process. Thus, sensing, control and information are inextricably tied. This has consequences in understanding the so-called ``signal-to-symbol barrier'' problem, as well as in the analysis and design of active sensing systems. I will show applications in vision-based control, navigation, 3-D reconstruction and rendering, as well as detection, localization, recognition and categorization of objects and scenes in live video.

Fritz Sommer, University of California, Berkeley

Information-Theory based policies for exploratory learning in sensorimotor loops

Over the last two decades great progress has been made in understanding how sensory representations are learned in the brain, driven by the principle of efficient coding. In contrast, we are still lacking theories of how motor output is guided for optimizing learning in closed sensor-motor loops. First, I will review foundational work that defined information gain and proposed it for guiding optimal experimental design and for driving learning in action-perception loops. Second, I will present results on exploratory learning of agents in unknown environments, whose actions optimize information gain within a multi-step time horizon. Finally, I will discuss information gain-driven learning in unbounded state spaces.

References:
1) DY Little, FT Sommer: Learning and exploration in action-perception loops. Front. Neural Circuits, 22 March 2013 | doi: 10.3389/fncir.2013.00037.
http://journal.frontiersin.org/Journal/10.3389/fncir.2013.00037/abstract
2) SA Mobin, JA Arnemann, FT Sommer: Information-based learning by agents in unbounded state spaces. Advances in Neural Information Processing Systems 27 (NIPS 2014).
http://papers.nips.cc/paper/5266-information-based-learning-by-agents-in-unbounded-state-spaces

Susanne Still, University of Hawaii

Optimal Information Processing: Dissipation and Irrelevant Information

Proposals for overarching biophysical principles have included both minimization of metabolic cost, as well as efficient processing of information. These are deeply related. In particular, optimization of thermodynamic efficiency often requires efficient information processing.

Here, I will derive an optimization principle for data representation from the thermodynamics of information engines, leading to some well known methods in information theory and machine learning. In particular, I show that minimizing the smallest achievable heat dissipation in an information engine leads to the Information Bottleneck method. Furthermore, I discuss how Shannon's information rate is proportional to the least effort necessary for data representation, and Shannon's channel capacity is related to maximum work potential.

Methods like rate-distortion theory and the Information Bottleneck have been adapted to model interactive learning behavior of agents with and without rewards. An interesting finding I wish to highlight is that behavioral policies that maximize predictive power have to balance between control and exploration. Exploration, in this framework, is then seen as a behavior driven by improving information gain, in contrast to policy randomization implemented by reinforcement learning's "Boltzmann-exploration".

Takashi Tanaka, KTH Royal Institute of Technology

LQG Control with Minimum Directed Information: A Semidefinite Programming Approach

Real-time decision-making procedures in general require continuous acquisition of information from the environment. In this talk, we revisit one of the most fundamental questions in real-time decision-making theory: what is the minimal information acquisition rate to achieve sequential decision-making with desired accuracy? We tackle this question using basic tools from control theory, information theory, and convex optimization theory. Specifically, we consider a Linear-Quadratic-Gaussian (LQG) control problem where Massey's directed information from the state sequence to the control sequence is taken into account. We show that the most "information-frugal" decision-making policy achieving desired LQG control performance admits an attractive three-stage separation structure comprised of (1) a linear sensor with additive Gaussian noise, (2) Kalman filter, and (3) a certainty equivalence controller. We also show that an optimal policy can be synthesized using a numerically efficient algorithm based on semidefinite programming (SDP).

Naftali Tishby, The Hebrew University of Jerusalem

The Synergy between Information and Control

The link between communication and control was obvious to Shannon, as evident from the conclusion of his 1959 paper on the duality between source and channel coding. Shannon noticed then the analogy between source coding and inference from past observations, and between channel coding and predictions of future events. Regrettably, the two fields grew apart, partly due to their different mathematical languages, the asymptotic (long delays) nature of Shannon's optimal coding versus the instantaneous nature of dynamical systems and optimal control. Yet, these two fundamental abstractions naturally merge when considering the joint-source-channel coding problem, with memory and feedback. The general distortion-cost tradeoff in communication can be compared, a-la Shannon, to the fundamental sensing-costs – actions-value tradeoff in optimal control and planning. The two theories begin to unify by considering optimal control under informational costs, as can be seen in several talks of this meeting. This unified view of information and control, when allowing adaptive channels and unspecified sources, connects directly with the fundamental problems of brain sciences and with the new computational methodologies of machine learning.

In this talk, I will sketch the intriguing analogies and synergy between communication and control and state some of the problems with their desired unification, as well as the main issues addressed in this meeting. I will then focus on the role of machine learning as a powerful framework for implementing this unification and applying it to the understanding of biological systems.

Daniel Wolpert, University of Cambridge

Sensorimotor decisions

The fields of decision making and sensorimotor control have developed in parallel although both require acting in real time on streams of noisy evidence. I will review our recent work showing the intimate interactions between decision making and sensorimotor control processes. This includes the relation between vacillation and changes of mind in decision making, the bidirectional flow of information between elements of decision formations and motor processes and the role of confidence in multi-stage decisions.