

PATRICK BAJARI/

Title: Estimating Dynamic Models with Strategic Interactions

Abstract: We describe two-step methods for estimating dynamic games under the assumption that behavior is consistent with Markov perfect equilibrium. In the first step, the policy functions and the law of motion for the state variables are estimated. In the second step, the remaining structural parameters are estimated using the optimality conditions for equilibrium. The second step estimator is a simple simulated minimum distance estimator. The algorithm applies to a broad class of models, including industry competition models with both discrete and continuous controls such as the Ericson and Pakes (1995) model. We discuss extensions of the method to estimating models where agents have heterogeneous beliefs, such as in the literature on agent based models.

NILS BERTSCHINGER (SECOND POSSIBILITY, WITH JUERGEN JOST AND ECKEHARD OLBRICH)/

Title: Autonomy and intentional action

Abstract: Strategic agents are described as acting according to internal incentives, e.g. motivations, utilities etc. For many natural as well as technical systems an intentional description is not readily available. Instead the system is described in terms of mechanisms and algorithms which generate its behavior. Here, we propose that autonomy, as a prerequisite of agency, can be identified in information theoretic terms. A system is called autonomous if it contains an internal degree of freedom which cannot be predicted from simply observing its interaction with the environment. Here, dependencies between the system and its environment are either attributed to the system, such as results of its actions, or considered as external influences. Only the latter should be included in the measure and reduce the system autonomy. The next step is then to interpret the internal structure of the system in terms of beliefs and goals. The system is then thought to act in order to achieve its goals. Even in the same environmental situation different behavior can be observed depending on the goal of the system. We propose that modal logic with modalities describing beliefs and goals of the system is a suitable framework to interpret the internal structure of autonomous agents. The logical framework allows for example to investigate how a strategic system takes into account beliefs about beliefs and goals of other systems.

STEFAN BIENIAWSKI:

Title: Exploring the Role of Health-Based Adaptation in Multi-Vehicle Missions Using Indoor Flight Experiments

Abstract: Significant investment has been made in the development of off-line systems for monitoring and predicting the condition and capability of aerospace systems. These are most typically used to reduce the operational costs of a system. A recent trend in aerospace is to include these technologies on-line and to utilize the provided information for real-time autonomous or semi-autonomous decision making. While forms of health-based adaptation are used commonly in critical functions, such as redundant flight control systems, as the scope is expanded - such as to the multiple vehicle level - new challenges and opportunities arise. For instance, the use of health based information in mission planning offers the opportunity to significantly enhance overall mission assurance. However, developing mission concepts, even at a simple level, requires coordination of multiple assets and determination of common interfaces suitable for heterogeneous fleets. For systems that are subject to real failures,

simulation offers the challenges of developing realistic scenarios and realistic health emulation. The approach taken, and reviewed in this presentation, has been to explore the domain using a sub-scale indoor flight test facility where real faults are common and manifest in different forms. The facility enables large numbers of flight hours and supports a wide range of vehicle types and component technologies. The approach allows exploration of a range of heterogeneous mission concepts providing better understanding of the interactions between individual vehicles as well as sub-systems within a vehicle. Of particular interest are persistent missions where faults are a key driver in the aggregate mission performance. Results of flight tests with several different sample missions will be presented. These missions range from non-cooperative to cooperative and include a range of tasks. The lessons learned and architecture are relevant for the broad range of aerospace systems. Future directions, such as the collaborative design of the core functions along with the health-based functions, will also be discussed.

MISHA CHERTKOV (REQUEST TO BE NEAR END OF SFI PORTION)/

Title: Smart Grid project at LANL and related challenges in learning and games.

SIMON DEDEO/

Title: Boltzmann Solution Concepts, epsilon Logic, and the Emergence of Timescales in an Animal Society

Abstract: Quantitative data on the behavior of animals in larger (N~50) groups allow for the detection and study of new phenomena that arise from the rational and perceptual capabilities of individuals acting in subgroup contexts. Here we report on three new approaches to a particular set of observations, of pigtailed macaques at the Yerkes Primate Research Center, that illuminate the complexity of group behavior in terms of game theory (Boltzmann Solution Concepts), noisy computational processes (epsilon-Logic), and the interaction of different environmental, social, physiological and cognitive mechanisms in the time domain (Lomb-Scargle periodogram analysis of timescales.)

HAIMONTI DUTTA/

Title: Modeling Load-Shifts and Failures in Feeder Networks: Progress and Challenges

Abstract:

JOE HALPERN/

Title: Distributed computing meets game theory: Fault tolerance and implementation with cheap talk

Abstract: Nash equilibrium is the most commonly-used notion of equilibrium in game theory. However, I argue that it does not have the robustness and fault tolerance properties that are important for applying it to distributed computing; in a precise sense, it does not tolerate "faulty" or "unexpected" behavior. I discuss notions of robust Nash equilibria, and show how and when a solution that achieves the desired robust equilibria using a mediator (trusted third party) can be implemented using what economists call "cheap talk", that is, by players just communicating among themselves. These results allow us to bring together over twenty years of work that has gone on largely independently in computer science and game theory. Joint work

with Ittai Abraham, Danny Dolev, and Rica Gonen.

JESSICA FLACK/

Title: Inductive game theory and collective conflict dynamics

DIRK HELBING/

Title: Self-Organization and Self-Optimization in Social and Traffic Systems

Abstract: The lecture presents models for social cooperation, pedestrians crowds, and traffic flows on freeways and urban road networks, discussing issues of self-organization and the outbreak and breakdown of coordination.

In particular, we propose new ways to support a fluid traffic operation. Through mechanism design, it is possible to specify local interactions between system elements such that it gives rise to a high performance of traffic systems. In other words, we show how one can create order through self-organization and reach high efficiency on a systemic level without the need for central control.

JUERGEN JOST/

Title: Some thoughts on the issue of rationality

Abstract: Rationality is a basic concept underlying economic and game theory. An agent in a game is assumed rational in the sense that she utilizes the best available strategy to maximize her utility, recognizing that her opponents are rational in the same sense. It can, however, be advantageous for a player in a game to act irrationally, in order to realize a better one among the possible Nash equilibria (example to be discussed: Quantal Response Equilibria) or to switch to a more advantageous game (example: persona games as higher level games). Moreover, the concept of rational expectations of economic theory may not be applicable in situations where the economic process not only causes the expectations of its participants, but is itself the result of the coordination of the expectations of its players (example: it can be rational to participate in an irrational bubble). This will lead us to the issue of mutual awareness between economic agents.

KALMANJE KRISHNAKUMAR/

Title: Decentralized control with human and (intelligent) artificial pilots - Benefits and Potential Pitfalls

Abstract: In this talk I will describe the current areas of research in making feedback control more intelligent such that it can be applied to both unmanned and piloted aerial vehicles. We will highlight some of the technologies studied with-in the NASA Aeronautics research portfolio and discuss the positive and (potentially) negative effects of such technologies in the context of decentralized control.

DAVID LESLIE (with Harriet Bass)/

Title: Controlled learning through taxation

Abstract: We present theoretical and simulated results on the control of game-theoretical learners by taxation. The method optimises the total tax revenue (or any other objective function) of the controller, while allowing the game players to learn. Changing of a tax rate is equivalent to changing the temperature parameter in a smooth best

response, such as a Boltzmann distribution. Hence the controller can move the players along a surface of quantal response equilibria in such a way as to improve the controller's reward. We prove that the controller will reach a local optimum of the long term average reward, and observe this fact in simulations.

RITCHIE LEE (with David Wolpert)/

Title: Using game theory to influence pilot behavior during near mid-air

Abstract: Traffic Alert and Collision Avoidance System (TCAS) is the current system deployed for warning pilots of possible mid-air collisions. Although pilots are trained to obey Resolution Advisories (RAs) during mid-air encounters, there is a wide variability in the way pilots actually respond. In fact, a recent study suggests that only 13% of pilot responses met the TCAS design assumptions in promptness and aggressiveness, with pilots acting in violation of the TCAS RA a whopping 24% of the time. The current TCAS system does not model this variability explicitly and only accounts for it indirectly via design buffers in threshold constants, and using extra-conservative rules.

The sources of pilot variability arise in three places: a) how the pilot perceives his/her environment, b) how pilots interacting in an encounter anticipate one's responses, and c) the's utility function. By combining concepts from Bayesian Networks and Game Theory into Network-Form, this work proposes a novel modeling methodology that enables the explicit modeling of the variability in the responses. In this framework, pilots are modeled as nodes in a Bayesian Network that defines their interaction with the environment in the context of the problem. Pilot behavior is modeled using Game Theory concepts such as Level-K Thinking and Sufficient Strategies, and the's response is ultimately decided by his/her utility function. This improved pilot model is a significant first step towards more accurate predictions of human behavior, opening the door to the design of improved RA systems. Furthermore, methods for optimizing any choice of performance metrics were investigated with promising results.

KEVIN LEYTON-BROWN (with Albert Xin Jiang, David R.M. Thompson, and Navin A.R. Bhat)/
Scaling Up Game Theory: Representation and Reasoning with Action Graph Games

<http://www.cs.ubc.ca/~kevinlb/papers/AGG.pdf> (main paper for this talk; to appear in GEB)

Abstract: Most work in game theory is analytic; it is less common to analyze a model's properties computationally. Key reasons for this are that game representation size tends to grow exponentially in the number of players--making all but the simplest games infeasible to write down--and that even when games can be represented, existing algorithms (e.g., for finding equilibria) tend to have worst-case performance exponential in the game's size. This talk describes Action-Graph Games (AGG), which make it possible to extend computational analysis to games that were previously far too large to consider. I will give an overview of our six-year effort developing AGGs, emphasizing the twin threads of representational compactness and computational tractability.

The first part of the talk will describe the core ideas of the AGG representation. AGGs are a fully-expressive, graph-based representation that can compactly express both strict and context-specific independencies in players' utility functions. I will illustrate the representation by describing several practical examples

of games that may be compactly represented as AGGs. The second part of the talk will examine algorithmic considerations. I'll describe a dynamic programming algorithm for computing a player's expected utility under a given mixed-strategy profile, which is tractable for bounded-in-degree AGGs. This algorithm can be leveraged to provide an exponential speedup in the computation of best response, Nash equilibrium, correlated equilibrium, and quantal response equilibrium. Second, I'll more briefly describe some current directions in our work on AGGs: a message-passing algorithm for computing pure-strategy Nash equilibria in symmetric AGGs, which is tractable for graphs with bounded treewidth; methods for performing computational analysis of real-world economic mechanisms; the extension of AGGs to both temporal and Bayesian-game settings; and the design of free software tools to make it easier for other researchers to use AGGs.

Our efforts in studying AGGs have tended to emphasize the analysis of existing systems (e.g., through various equilibrium concepts) rather than the design and control of novel systems. I'll be interested in speaking with other workshop attendees both during this talk and afterwards about how we might apply our techniques to addressing control problems.

MICHAEL LITTMAN (joint work with Michael Wunder and Monica Babes)/
Title: Classes of Multiagent Q-learning Dynamics with epsilon-greedy Exploration

Abstract: The Q-learning reinforcement-learning algorithm is known to converge to optimal behavior in the limit in single-agent environments given sufficient exploration. The same algorithm has been applied, with some success, in multiagent environments, where traditional analysis techniques break down. Using dynamical systems methods, we derived and studied an idealization of Q-learning in 2-player 2-action repeated general-sum games. In particular, we address the discontinuous case of epsilon-greedy exploration and use it as a proxy for value-based algorithms to highlight a contrast with existing results in policy search. Analogously to previous results for gradient ascent algorithms, we provide a complete catalog of the convergence behavior of the epsilon-greedy Q-learning algorithm by introducing new subclasses of these games. We identify two subclasses of Prisoner's Dilemma-like games where the application of Q-learning with epsilon-greedy exploration results in higher-than-Nash payoffs for a range of initial conditions.

MATTEO MARSILI (with A. Kirman, N. Hanaki and P. Pin)/
Title: Ownership by luck

Abstract Consider a generic situation where a population of agents asynchronously accesses a number of resources. Usage of resources is exclusive: if an agent is using a resource, other agents cannot use it. Examples include searching for parking, establishing colonies and animals trying to establish a territory or a position in pecking order.

Nash equilibria can be of two types: Symmetric, when each agent adopts the same strategy, and asymmetric, when different agents play differently. When asymmetric outcomes prevail, some agents may turn out to occupy more frequently the best resources, as if they were lucky, or if they had property rights on those resources.

When agents rank resources differently, asymmetric outcomes are expected. When resources are equivalent, the problem becomes one of coordination. Again asymmetric outcomes are (evolutionarily) selected. When there is an objective ranking of resources (i.e. everybody

regards resource a as better than resource b) the situation is more complex, as incentives and the cost of mis-coordination compete.

I discuss, in simple models, the transition from symmetric to asymmetric states, how it materializes and its determinants.

FRANS OLIEHOEK/

Title: Exploiting structure in collaborative games with private information.

Abstract: This talk focuses on collaborative decision making under uncertainty: settings in which agents share the same payoff function, but each agent may have a different partial view of its environment. One shot interactions in such settings can be modeled by collaborative Bayesian games (CBGs), in which each agent has a particular type that defines the private information it has about the environment.

There are two main issues that prevent the CBG framework from scaling up: finding a solution (a Pareto optimal Nash equilibrium) is NP-hard, and the representation itself scales exponentially with the number of agents. These problems have been addressed independently of each other: graphical games exploit structure of independence between agents to allow for the representation of many agents, other recent work exploits the structure between types to find solutions more efficiently.

In this work, we propose the collaborative graphical BG (CGBG) as a model that extends the graphical game formulation to CBGs and propose a solution method that exploits both types of structure. We show how 1) a CGBG corresponds to a factor graph that represents both types of structure in a uniform way, and 2) the problem can be approximately solved by running message passing over this factor graph. Finally, we consider the impact of our results in sequential settings modeled by decentralized partially observable Markov decision processes (Dec-POMDPs). We show that CGBGs and their efficient solution allows for the approximate solution of Dec-POMDPs with hundreds of agents.

OZAN CANDOGAN (with Asu Ozdaglar, Ishai Menache, and Pablo Parrilo)/

Title: Flow Representations of Games: Near Potential Games and Dynamics

Abstract: Despite much interest in using game theoretic models for the analysis of resource allocation problems in multi-agent networked systems, most of the existing works focus on static equilibrium analysis without establishing how an equilibrium can be reached dynamically. In the theory of games, natural distributed dynamics reach an equilibrium only for restrictive classes of games; potential games is an example. These considerations lead to a natural and important question: can we have a systematic approach to analyze dynamic properties of natural update schemes for general games?

Motivated by this question, this talk presents a new approach for the analysis of games, which involves viewing preferences of agents over the strategy profiles as flows on a graph. Using tools from the theory of graph flows (which are combinatorial analogues of those for continuous vector fields), this allows investigating topological properties of preferences. In particular, we use a flow-decomposition theorem, Helmholtz Decomposition theorem, to show that any finite strategic form game can be written as the direct sum of a potential game, a harmonic game, and a nonstrategic part. Hence, this decomposition leads to a new class of games, "harmonic games", with well-understood equilibrium and dynamic properties. Moreover, this approach allows projecting an arbitrary game onto the space of potential games (or harmonic games) using convex optimization and

exploit the relation between the two games to analyze the static and dynamic equilibrium properties of the original game. The second part of the talk uses this idea to study a non-cooperative power control game and characterize the system optimality properties along dynamic trajectories of natural user update schemes for this game.

ECKEHARD OLBRICH (with N. Bertschinger, A. Kabalak, J.Jost)/

Title: Communication in systems of interacting strategic agents

Abstract: An essential part of human cooperation is communication. Therefore it would be natural to ask for the role of communication in systems of interacting strategic agents, either artificial or mixed artificial and human. The first problem is to define communication in such a setting. In particular one can ask, how communication can be distinguished from pure interaction. We propose a concept of communication that distinguishes different levels of complexity starting from the simple interaction between two systems that generates mutual information between the system states that can be encountered already on the level of physical systems and will end with a notion of communication that incorporates specific aspects of human communication as it is formulated in the openness condition by Grice. Any lower level is a necessary, but not sufficient condition for the next higher level. Moreover, the different levels correspond to different descriptions. While the lowest level corresponds to the physical description as a dynamical system the higher levels require notions such as 'belief', 'intention' or 'beliefs about intentions', for instance by using modal logic. Moreover, the occurrence of higher levels of communication should correspond to specific properties on the level of the physical description. Therefore a translation between the different levels of description should be helpful both for designing artificial and understanding artificial or natural systems of interacting strategic agents.

BRIAN ROGERS/

Title: Emergence of Cooperation in Anonymous Social Networks through Social Capital

Abstract: We study the emergence of cooperation in dynamic, anonymous social networks, such as in online communities. We examine prisoner's dilemma played under a social matching protocol, where individuals form random links to partners with whom they can interact. Cooperation results in mutual benefits, whereas defection results in a high short-term gain. Moreover, an agent that defects can escape reciprocity by virtue of anonymity: it is always possible for an agent to abandon his history and re-enter the network as a new user. We find that cooperation is sustainable at equilibrium in such a model. Indeed, cooperation allows an individual to interact with an increasing number of other cooperators, resulting in the formation of a type of social capital. This process arises endogenously, without the need for potentially harmful social enforcement rules. Additionally, for a rich class of parameter settings, our model predicts a stable coexistence of cooperating and defecting agents at equilibrium.

SUJAY SANGHAVI/

Title: Belief Propagation for Networks

Abstract: Belief Propagation (BP) is a message-passing algorithm developed for large-scale estimation and inference problems in statistical physics and machine learning. In this talk we overview recent research that

shows its effectiveness in a very different application domain: distributed resource allocation in networks. In particular, we show that BP, and related algorithms, have some very appealing properties in these settings, and also highlight the challenges that prevent BP to be used "out of the box", and our modifications to circumvent the same.

Time permitting, we will draw connections between BP and popular auction mechanisms, like the Vickrey-Clarke-Groves (VCG) auction, in distributed settings. In particular, we show the correspondence between BP updates and a natural myopic bid update rule for VCG auctions.

DUSAN STIPANOVIC/

Title: Accomplishing Multiple Objectives by Multiple Agents using Convergent Approximations of the Min and Max Functions

Abstract: In this talk, we will present an approach based on convergent and continuously differentiable approximations of the min and max functions to design strategies for agents aiming to accomplish multiple objectives. The conditions that guarantee an accomplishment of multiple objectives are based on differential inequalities and minimal and maximal solutions of differential equations. We associate an objective function to each objective and construct agents' goal functions using approximations of the min and max functions acting as logical "and" and "or" functions. Then we use differential inequalities and the comparison principle to establish conditions guaranteeing that the objectives will be accomplished. By doing so, we bypass solving Hamilton-Jacobi-Bellman-Isaacs partial differential equations and in some relevant cases can even provide closed-form solutions for agents' strategies.

NAFTALI TISHBY (WITH JONATHAN RUBIN AND OHAD SHAMIR)/

Title: Robust optimal control by trading future information and value

Abstract: One of the most striking characterizations of life is the ability to efficiently extract information - through sensory perception, and exploit it - through behavior. There is a growing empirical evidence that information seeking is as important for optimal behavior as reward seeking. Yet our basic algorithms for describing planning and behavior, in particular reinforcement learning (RL), so far ignored this component. In this talk I will describe new extensions of reinforcement learning that combine information seeking and reward seeking behaviors in a principle optimal way. I will argue that Shannon's information measures provide the only consistent way for trading information with expected future reward and show how the two can be naturally combined in the frameworks of Markov-Decision-Processes (MDP) and Dynamic Programming (DP). This new framework unifies techniques from information theory (like the Huffman source coding algorithm) with methods of optimal control (like the Bellman equation). We show that the resulting optimization problem has a unique global minimum and convergence (even that it lacks convexity). Moreover, the tradeoff between information and value is shown to be robust to fluctuations in the reward values by using the PAC-Bayes generalization bound, providing another interesting justification to its biological relevance.

BRENDAN TRACEY (with David Wolpert and Juan Alonso)/

Title: Using supervised learning to improve Monte Carlo integral estimation

Abstract: Monte Carlo (MC) techniques are used to estimate integrals of a function using randomly generated samples of the function. While MC techniques have proven to be one of the most powerful tools of science and engineering, they often suffer from high variance and slow convergence.

In this talk we present Stacked Monte Carlo (StackMC), a new method for postprocessing a given set of MC samples to improve the associated integral estimate. In theory stackMC reduces the variance of any type of Monte Carlo integral estimate (simple sampling, importance sampling, quasi-Monte Carlo, MCMC, etc.) without adding bias. We report an extensive set of experiments confirming that the stackMC estimate of an integral is more accurate than both the associated pre-processing Monte Carlo estimate and an estimate based on a functional fit to the MC samples. These experiments run over a wide variety of integration spaces, numbers of sample points, dimensions, and fitting functions.

JAMES WRIGHT/

Title: Beyond Equilibrium: Predicting Human Behavior in Normal Form Games

Abstract: It is standard in multiagent settings to assume that agents will adopt Nash equilibrium strategies. However, studies in experimental economics demonstrate that Nash equilibrium is a poor description of human players' actual behavior. In this study, we consider a range of widely studied models from behavioral game theory. For what we believe is the first time, we evaluate each of these models in a meta-analysis, taking as our data set large-scale and publicly-available experimental data from the literature. We then propose a modified model that we believe is more suitable for practical prediction of human behavior.

PEYTON YOUNG (WITH BARY S. R. PRADELSKI)

Title: Efficiency and Equilibrium in Trial and Error Learning

Abstract. In trial and error learning, agents experiment with new strategies and adopt them with a probability that depends on their realized payoffs. Such rules are completely uncoupled, that is, each agent's behaviour depends only on his own realized payoffs and not on the payoffs or actions of anyone else. We show that there is a simple version of trial and error learning that selects a Pareto optimal equilibrium whenever a pure equilibrium exists, no matter how large or how complex the game may be. In games where a pure equilibrium does not exist, the long-run likelihood of every disequilibrium state is determined by a weighted combination of two factors: the total payoff to all agents in that state, and the maximum payoff gain that would result from a unilateral deviation by some agent. This welfare/stability trade-off criterion provides a novel framework for analyzing the selection of disequilibrium as well as equilibrium states in finite n-person games.

JOEL WATSON/

TBA

DAVID WOLPERT (WITH JAMES BONO)/

Title: Solution Concepts that are Distributions over Profiles rather than Sets of Profiles

Abstract: Conventionally, game theory predicts that the mixed strategy profile of players in a particular noncooperative game will fall within some set determined by the game, e.g., the set of Nash equilibria of that game. Relative probabilities of strategy profiles in that set are unspecified, and all profiles not in the set are implicitly assigned probability zero. However the axioms underlying Bayesian rationality tell us to predict the state of a system using a probability density over the set of all possible states, not using a subset of all possible states. So when the "set of all possible states" is the set of mixed strategy profiles of a game, Bayesian rationality tells us to use a density over the set of all profiles, not a subset of such profiles. Via standard Bayesian decision theory, such a density provides a best single prediction of the profile of any noncooperative game, i.e., a universal refinement. In addition, regulators can use such a density to make Bayes optimal choices of a mechanism, thereby fully adhering to Savage's axioms. In particular, they can do this in strategic situations where conventional mechanism design cannot provide advice. We illustrate all of this on a Cournot duopoly game.

COREY IPPOLITO/
No presentation.

LENKA/
No presentation.