David McDonald (University of Ottawa)

Yaglom Limits - **tutorial**

As Keynes said, "In the long run we are all dead". Does it really make sense to study steady state when in fact all processes are evanescente. Shouldn't we rather study processes conditional on the fact they are still alive.

Yaglom limits are the limiting distribution of an absorbing Markov chain X_n with kernel K on state space S given it has survived up to time n; i.e., the limit of $K^n(x, y)/K^n(x, S) \to \pi(y)$. We discuss the enormous literature related to this limit as well as some recent work showing π can depend on the starting point x.

Wei Sun (Concordia University)

Hunt's hypothesis (H) and Getoor's conjecture

A Markov process X is said to satisfy Hunt's hypothesis (H) if every semi-polar set is polar. Roughly speaking, this means that if a set A cannot be immediately hit by X for any starting point, then A will never be hit by X. Which Lévy processes satisfy Hunt's hypothesis (H) is a long-standing open problem in probabilistic potential theory. In this talk, I will summarize the results that we have obtained for this problem in recent years. In particular, I will present theorems and examples on the validity of (H) for one-dimensional Lévy processes and the sum of two independent Lévy processes. This talk is based on joint papers with Zechun Hu.

Jhelum Chakravorty (McGill University)

Fundamental Limits of Remote Estimation

In many applications of real-time communication systems such as sensor networks, smart grids, transportation networks, internet of things etc., there are sequential transmissions of data. In such scenarios, a transmitter (sensor) observes a Markov source realization and sends its observation to a remotely placed receiver, connected by a communication channel. The transmitter is a battery powered device that transmits over a wireless packet-switched network. Switching the radio on is significantly more expensive compared to the sensing and the duration for which the transmitter is kept on. Hence the transmission is intermittent but whenever there is a transmission, the entire source symbol gets transmitted. Whenever there is no transmission, the receiver has to estimate the source realization. In this talk, we characterize fundamental trade-offs between the estimation error (or distortion) and the cost of transmission in such systems.

In particular, we consider the case when the source is a first-order autoregressive process and the transmissions take place over a Markov erasure (or Gilbert-Elliott) channel. We characterize the following two fundamental trade-offs: (a) for costly transmission, what is the minimum expected distortion plus transmission cost and (b) when there is a constraint in the expected number of transmissions, what is the minimum expected distortion. For both cases, we establish the structure of optimal strategies, characterize the optimal performances and provide analytical as well as numerical methods to compute the solution. This is joint work with Jayakumar Subramanian and Aditya Mahajan.

Peter Caines (McGill University)

Mean Field Game Theory for Partially Observed Systems with Applications to Execution Problems in Finance

Mean field game theory for partially observed systems (PO-MFG theory) studies systems with a population of asymptotically negligible minor agents and, possibly, a major agent, where, in the general case, all minor agents partially observe their own and the major agent's state, and the major agent partially observes its own state.

The optimal execution problem in financial markets may be formulated in terms of a population of high frequency traders (minor agents) together with an institutional (major) investor, all of whom are (i) coupled through the mean field inherent in the market, and (ii) have partial observations on the their individual states and on the major agent's inventory. Furthermore, each agent is assumed to have a utility function reflecting the agents? interest in maximizing their individual terminal wealth while avoiding both large execution prices and high trading rates. An application of PO-MFG theory developed for linear stochastic systems with quadratic utility functions gives the best rate of trading for each agent to maximize its utility function in the ε-Nash equilibrium sense. Computational examples will be presented together with extensions of PO-MFG theory to financial stopping problems using hybrid stochastic optimal control theory.

Work with D. Firoozi, A.Kizilkale, A. Pakniyat and N. Sen.

Geneviève Gauthier (HEC Montreal)

Extracting Latent States from High Frequency Option Prices

We propose the realized option variance as a new observable variable to integrate high frequency option prices in the inference of option pricing models. Using simulation and empirical studies, this paper documents the incremental information offered by this realized measure. Our empirical results show that the information contained in the realized option variance improves the inference of model variables such as the instantaneous variance and variance jumps of the S&P 500 index. Parameter estimates indicate that the risk premium breakdown between jump and diffusive risks is affected by the omission of this information.

Ali Kara (Queen's University)

Continuity and Robustness to Incorrect Priors for Partially Observed Stochastic Control

We will look at the continuity properties of single and multi-stage stochastic control problems with respect to initial probability measures and applications of these results to the study of robustness of optimal control policies applied to systems with incomplete or incorrect probabilistic models. It is shown that for single and multi-stage optimal cost problems, continuity and robustness cannot be established under weak and setwise convergences in general, but that the optimal cost is continuous in the priors under the topology of total variation under mild conditions. By imposing further assumptions on the measurement models, we also show that the optimal cost can be made continuous under weak convergence of priors. Using these continuity results we show that if an optimal control policy is applied to a prior model Q, and if Q is close to the true model P, then the application of the incorrect optimal policy to the true model leads to a loss that is continuous in the distance between Q and P under total variation, and under further conditions, in the Wasserstein distance and weak convergence, leading to a form of robustness. Implications of these results in empirical learning for control will be presented, where weak convergence of empirical measures typically occurs but stronger notions of convergence do not.

Joint work with Graeme Baker and Serdar Yuksel.

Ravi Mazumdar (University of Waterloo)

Insensitivity of the Mean Field of Loss Systems under Randomized SQ(d) Algorithms

In many applications such as cloud computing, managing server farm resources etc. an incoming task or job has to be matched with an appropriate server in order to minimise the latency or blocking associated with the processing. Ideally the best choice would be to match a job to the fastest available server. However when there are thousands of servers requiring the information on all server tasks is an overkill.

Pioneered in the 1990's the idea of randomized sampling of a few servers was proposed by Vvedenskaya and Dobrushin in Russia and Mitzmenmacher in the US and popularized as the "power of two" schemes which basically means that sampling two servers randomly and sending the job to the "better" server (i.e. with the shortest queue, or most resources) provides most of the benefits of sampling all the servers.

In the talk I will discuss multi-server loss models under power-of-d routing scheme when service time distributions are general with finite mean. Previous works on these models assume that the service times are exponentially distributed and insensitivity was suggested through simulations. Showing insensitivity to service time distributions has remained an open problem. The difficulty is that for general service times the underlying Markovian model is more complex. Using a measure valued process approach we first derive the mean field equation (MFE) for the empirical measure. The MFE is now characterized by a PDE whose stationary point coincides with the fixed point in the case with exponential service times. This establishes the insensitivity of the fixed point. The techniques developed in this paper are applicable to study mean field limits for Markov processes on general state spaces and insensitivity properties of other queueing models.

Ryan Kinnear (University of Waterloo)

Learning Granger Causality Graphs: Causal Inference for Time Series Data

Identifying causal (rather than merely correlative) relationships in dynamic systems is a difficult task, particularly if it is not feasible to perform controlled experiments. Granger's notion of causality was developed first in economics beginning in the 1960s and can be used to form a network of plausible causal relations given only the opportunity to observe the system. This method is applied, for example, in neuro-imaging to identify relationships amongst brain regions, and in biostatistics to explore gene regulatory networks. In this talk, we provide an overview of the notion of Granger Causality, some applications, and methods for learning Granger Causality Networks in practice.

Aditya Mahajan (McGill University)

Decentralized Kalman Filtering

We consider a decentralized estimation problem in which agents observe the state of a dynamical system with noise. Agents are connected over a communication graph. At each time, each agent transmits all its information to its neighbours and generates an estimate of the state of the system. The per-step estimation error depends on how close the estimates are to the true state of the system and *how close the estimates of the agents are to their neighbours* (according to a cost graph, which may be different from the communication graph). We show that this problem can be mapped to a static team and hence, the optimal estimates are linear. We also develop present a method to compute the estimates recursively. Joint work with Mohammad Afshari.

Serdar Yuksel (Queen's University)

Strategic Measures Approach to Decentralized Stochastic Control: Structural, Existence and Approximation Results

This talk is concerned with decentralized stochastic control (or dynamic team) problems and their optimal solutions. After a brief review on decentralized stochastic control, strategic measures for such problems will be introduced; these are the probability measures induced on the space of measurement and action sequences by admissible decentralized control policies which satisfy various conditional independence properties. Conditions ensuring the compactness of sets of strategic measures will be established for both static and dynamic teams. These will lead to existence results on optimal solutions for both static and dynamic teams. The results are applicable to teams which are either static or static reducible, as well as teams which are classical. Properties such as convexity and Borel measurability of the sets of such measures will be studied; it will be shown that measures induced by deterministic policies form the extreme points of a properly expanded set of strategic measures, thus establishing the optimality of such policies. It will be shown that such sets are Borel, leading to positive implications for a general form of dynamic programming for sequential dynamic teams. Finally, through a proper approximation of the sets of strategic measures by those induced with quantization of measurement and action spaces, asymptotic optimality of finite model representations for a large class of dynamic team problems will be established. These lead to asymptotic optimality of quantized control policies. The celebrated counterexample of Witsenhausen will be discussed throughout the talk to illustrate the salient aspects of information structures in decentralized stochastic control, and demonstrate the existence and approximation results.

Roland Malhame (Ecole Poly. de Montreal)

Min-LQG Games and Collective Discrete Choice Problems

We introduce a novel class of finite horizon linear quadratic Gaussian games involving distinct potential finite destination states, interpreted as discrete choices under social pressure. The model provides stylized interpretations of opinion swings in elections, the dynamics of discrete societal choices, as well as a framework for achieving communication constrained group decision making in micro-robotic based exploration.

Two distinct cases are considered: (i) The zero noise or 'deterministic' case where agents are initially randomly distributed over their range space; (ii) The fully stochastic case. Under mild technical conditions, the existence of ε -Nash equilibria is established in both cases although these equilibria may in general be multiple. The corresponding agent control strategies are of a decentralized nature and are characterized in each case by the fixed points of a specific finite dimensional operator. Individual agent destination choices are fixed at the outset in case (i), while by contrast, their probability distribution evolves randomly along trajectories in case (ii), with a deterministic limit for the complete population as the latter grows to infinity. This is joint work with Rabih Salhab and Jerome Le Ny.

Alex Shestopaloff (University of Toronto)

Sampling Latent States for High-dimensional Non-linear State Space Models with the Embedded HMM Method

We propose a new scheme for selecting pool states for the embedded Hidden Markov Model (HMM) Markov Chain Monte Carlo (MCMC) method. This new scheme allows the embedded HMM method to be used for efficient sampling of state sequences in state-space models where the state can be high-dimensional. Previously, embedded HMM methods were only applied to models with a one-dimensional state space. We demonstrate that using our proposed pool state selection scheme, an embedded HMM sampler can have similar performance to a well-tuned sampler that uses a combination of Particle Gibbs with Backward Sampling (PGBS) and Metropolis updates. The scaling to higher dimensions is made possible by selecting pool states locally near the current value of the state sequence. The proposed pool state selection scheme also allows each iteration of the embedded HMM sampler to take time linear in the number of the pool states, as opposed to quadratic as in the original embedded HMM sampler. We also consider a model with a multimodal posterior, and show how a technique we term "mirroring" can be used to efficiently move between the modes. Joint work with Radford M. Neal.

Vladimir Vinogradov (Ohio University)

On Branching Particle Systems and Galton-Watson Processes with Sibuya-Type Branching Mechanism

We prove that a continuous-time branching particle system which starts from a discrete stable number of particles and is characterized by a scaled Sibuya branching mechanism must have evolved from a Poisson field. We provide a local version of this backward evolution result for such a system of Neveu-type Markov branching processes and derive other local limit theorems for its marginals. We consider the dual subcritical Galton-Watson process constructed starting from a discrete-time Galton-Watson process with a scaled Sibuya branching mechanism conditioned by extinction. For the dual process, we derive the probability-generating function of the law of the total progeny. We establish local approximation for the long-time limit of a Poisson field of such independent "dual" processes.

Shui Feng (McMaster University)

Asymptotic Results of Two-Parameter Dirichlet Process

The Dirichlet process is a random measure that has been widely used in Bayesian statistics and population genetics. The connection to exchangeable random partitions led to the development of its two-parameter generalization. In this talk we will discuss several asymptotic results associated with different parametric domains.

Xiaowen Zhou (Concordia University)

A Continuous-state Nonlinear Branching Process

A continuous-state branching process can be identified as the unique nonnegative solution to a SDE driven by a Brownian motion and a compensated Poisson random measure; see (J. Bertoin and J.-F. Le Gall (2006): Stochastic flows associated to coalescent processes III: Limit theorems. *Illinois J. Math.* **50**, 147–181.) and (D. Dawson and Z. Li (2012): Stochastic equations, flows and measure-valued processes. *Ann. Probab.* **40**, 813–857). By adapting this SDE, we can introduce a continuous-state branching process with nonlinear branching mechanism. Intuitively, the solution to the modified SDE is a branching process with branching rates depending on the current population size.

Using a martingale approach, we study its survival/extinction behaviours and find respective sufficient conditions on the branching parameters under which the process either survives with probability one or dies out with a positive probability. Similarly, we can also discuss the explosion behaviours for the continuous-state nonlinear branching process. We will show that those conditions are quite sharp.

Jun Yang (University of Toronto)

Complexity Bounds for Markov Chain Monte Carlo

This talk considers whether MCMC quantitative convergence bounds can be translated into complexity bounds. We prove that a certain realistic Gibbs sampler algorithm converges in constant number of iterations. Our proof uses a new general method of establishing a generalized geometric drift condition defined in a subset of the state space. The subset is chosen to rule out some "bad" states which have poor drift property when the dimension gets large. Using the new general method, the obtained quantitative bounds for the Gibbs sampler algorithm can be translated to tight complexity bounds in high-dimensional setting. It is our hope that the new general approach can be employed in many other specific examples to obtain complexity bounds for high-dimensional Markov chains.

Opher Baron (University of Toronto)

Queueing and Markov Chain Decomposition (QMCD),
the Single Stage Subsystems Case: Motivation and Examples - **tutorial**

We introduce Queueing and Markov Chain Decomposition (QMCD) for the exact analysis of Queues and Markov Chains. QMCD includes four steps: Decompose the system to smaller subsystems, Tie the subsystems together - transition rates and other effects, Solve each subsystem while considering the relevant effects, and Normalize the solution. We focus on applications of QMCD where the subsystems are single stage systems, such as the $M_n/G_n/1$ with state-dependent arrivals and services. We also discuss QMCD with more complicated subsystems such as for the $M/M/C$ with preemptive priority and different service rates, and a threshold based allocation of Blood units to patients. QMCD allows us to provide exact analysis of these, previously well studied but unsolved, models. We will discuss current and future research.

Based upon Joint work with Abouee-Mehrizi, H., Barron, Y., Berman, O., Balcioglu, B., Economou, A., Manou, A., Sarhangian, V., Scheller Wolf, A., Wang J.

Javad Tavakoli (UBC)

Numerical Methods to Deal with GI/G/1 Queues when Inter-arrival Times and/or Service Times have Geometric Tails

We discuss a number of methods to find the distributions of the waiting time, the idle time, and the length of discrete-time of $GI/G/1$ queue when inter-arrival or service time distributions have geometric tails. First, we find the waiting-time and idle time distributions by a modification of an algorithm suggested earlier by W. Grassmann and J.L.Jain. Next, we present three methods for finding the distribution of the number of all elements in the system. In the first method, we formulate a Markov chain with three state variables, the length of the line, the time since the last arrival, and, if the server is busy, the time since service started. The next method uses a Markov chain embedded at the time service has started. Finally, we show how the distribution of the number in the system can be found from the waiting-time distribution. Numerical and theoretical arguments show that this last method is the most efficient one, often by several orders of magnitude.

Haosui Duanmu (University of Toronto)

Nonstandard Analysis and its Application to Markov Processes

I will start by giving a short introduction on Nonstandard Analysis and Hyperfinite Probability Theory. I will then give then describe the basic idea of hyperfinite representations of probability spaces. A classical example of hyperfinite representation on unit interval [0,1] with Lebesgue measure will be discussed. Then I will introduce the notion of Hyperfinite Markov processes and prove the Markov chain ergodic theorem for hyperfinite Markov processes. We can construct a hyperfinite Markov process for continuous-time general state space Markov processes satisfying certain regularity conditions. The ergodicity of hyperfinite Markov processes imply the ergodicity of the original standard Markov processes.

Jalal Khamse Ashari (Carleton University)

Multi-resource fair allocation/scheduling for heterogeneous servers

Users of cloud computing platforms pose different types of demands for multiple resources on servers (physical or virtual machines). Besides differences in their resource capacities, servers may be additionally heterogeneous in their ability to service users -

certain users' tasks may only be serviced by a subset of the servers. We identify important shortcomings in existing multi-resource fair allocation mechanisms - Dominant Resource Fairness (DRF) and its follow up work - when used in such environments. We develop a new fair allocation mechanism, called Per-Server Dominant-Share Fairness (PS-DSF), which we show offers several desirable properties that DRF is able to offer for a single server.

We will discuss how our proposed per-server fairness criteria can be used to develop a fully distributed multi-resource allocation/scheduling algorithm.