Universiteit Leiden
Mathematical Institute

Applied Probability Conference

Current Directions in Applied Probability & Data Science

April 21-22, 2023

via Zoom (Time Zone NY, USA)

In honor of Michael Katehakis' 70th birthday and Isaac Sonin's 80th birthday

Time Schedule

	21-Apr			22-Apr	
	NY time			NY time	
Welcome	8.00-8.30				
Session I			Session I		
I	8.30-8.55	Jerzy Filar	I	8.30-8.55	Panagiotis Karras
II	8.55-9.20	Katsunori Ano	II	8.55-9.20	Apostolos Burnetas
III	9.20-9.45	Konstantin Avrachenkov	III	9.20-9.45	Matthew Sobel
Session II			Session II		
IV	9.45-10.10	Benjamin Melamed	IV	9.45-10.10	Chris Tsokos
V	10.10-10.35	Peter Jacko	V	10.10-10.35	Eugene Feinberg
VI	10.35-11.00	Krzysztof Szajowski	VI	10.35-11.00	Tze Lai
Break	11.00-11.30		Break	11.00-11.30	
Session III			Session III		
VII	11.30-11.55	Endre Boros	VII	11.30-11.55	Andrzej Ruszczynski
VIII	11.55-12.20	Saul Jacka	VIII	11.55-12.20	Xinxin Hu
IX	12.20-12.45	Stella Kapodistria	IX	12.20-12.45	Farid Alizadeh
Session IV			Session IV		
X	12.45-13.10	Benjamin Van Roy	X	12.45-13.10	Ray Cao
XI	13.10-13.35	Rhonda Righter	XI	13.10-13.35	Jian Yang
XII	13.35-14.00	Pavlo Kasyanov	XII	13.35-14.00	Jaideep Vaidya
Break	14.00-14.30		Break	14.00-14.30	
Session V			Session V		
XIII	14.30-14.55	Dimitri Bertsekas	XIII	14.30-14.55	Sadegh Talebi
XIV	14.55-15.20	Jiaqiao Hu	XIV	14.55-15.20	Dimitris Metaxas
XV	15.20-15.45	Sergei Schreider	XV	15.20-15.45	Jim (Junmin) Shi
XVI	15.45-16.10	Csaba Szepesvari	XVI	15.45-16.10	Karti Puranam
Break	16.10-16.40		Break	16.10-16.40	
Session VI			Session VI		
XVII	16.40-17.05	John P.A. Ioannidis	XVII	16.40-17.05	Victor de la Pena
XVIII	17.05-17.30	Peter Glynn	XVIII	17.05-17.30	Davit Sargsyan
XIX	17.30-17.55	Yao Zhao	XIX	17.30-17.55	Yajie Duan
XX	17.55-18.20	Sheldon Ross	XX	17.55-18.20	Debopriya Ghosh
Break	18.20-18.45		Break	18.20-18.45	
Open session	18.45-19.45		Open session	18.45-19.45	

Conference Moderator: Odysseas Kanavetas, Leiden University, The Netherlands.

Friday April 21, 2023.

Back to the time schedule.

Opening Remarks. 8.00 - 8.30 AM

Floske Spieksma, Leiden University, The Netherlands.
Lei Lei, Rutgers Business School, USA.
Apostolos Burnetas, National and Kapodistrian University of Athens, Greece.
Matthew Sobel, Case Western Reserve University, USA.
Odysseas Kanavetas, Leiden University, The Netherlands.
Back to the time schedule.

Session 1. 8.30 - 9.45 AM. Konstantin Avrachenkov, INRIA Sophia Antipolis, Chair.

Hidden Equations of Threshold Risk.

Jerzy Filar, The University of Queensland, Australia.

We consider the problem of parametric sensitivity of a particular characterization of risk, with respect to a threshold parameter δ. Such threshold risk is modeled as the probability of a δ−perturbed function of a random variable falling below 0. We demonstrate that for polynomial and rational functions of that random variable there exist at most finitely many risk critical points. The latter are those special values of the threshold parameter for which rate of change of risk is unbounded as δ approaches them. Under weak conditions, we characterize candidates for risk critical points as zeroes of either the discriminant of a relevant δ−perturbed polynomial, or of its leading coefficient, or both. Thus, the equations that need to be solved are themselves polynomial equations in δ that exploit the algebraic properties of the underlying polynomial or rational functions. We also extend the analysis to multidimensional case of δ−perturbed polynomials, where δ is a vector of parameters. This is joint work with V. Ejov and Z. Qiao.

Keywords: Threshold Risk, Discriminant, Varieties.

Back to the time schedule.

On Relation Between Probability of Win and Threshold Values of Optimal Stopping Rule for Multiple Stopping Secretary Problem.

Katsunori Ano, Office K.A Co., Ltd, Japan.

For secretary problem, it is known that the asymptotic probability of win is 1/e and the asymptotic threshold value of optimal stopping rule is also 1/e. For multiple stopping secretary problem, it is shown in Matsui and Ano (2016) that the asymptotic probability of win is sum of the asymptotic threshold values of optimal multiple stopping rule. We consider multiple stopping secretary problem with random number of objects. Under PS condition (Presman and Sonin (1972)) on the distribution of random number, optimal stopping rule is shown to be threshold type. Our conjecture is that under PS condition, the asymptotic probability of win is also sum of the asymptotic threshold values of optimal multiple stopping rule, which has not been proved yet, but it is an interesting relation.

Keywords: optimal stopping, odds problem, secretary problem, multiple stopping, probability of win.

Back to the time schedule.

Markov Chains with Restart.

Konstantin Avrachenkov, INRIA Sophia Antipolis, France.

Motivated by numerous applications in telecommunications, computer science and physics, we consider Markov chains with restart. At each time step, such a chain either with a positive probability restarts from a given distribution, or with the complementary probability continues according to a Markov transition kernel. We review the properties of Markov chains with restart, in particular, the characterization of the hitting times. If time permits, we discuss the applications of Markov chains with restart to network centrality indices and network sampling.

Keywords: Markov chains, Restart, Hitting times, Network centralities, Network sampling.

Back to the time schedule.

Session 2. 9.45 - 11.00 AM. Peter Jacko, Berry Consultants and Lancaster University, Chair.

Valuation of Supply Chains: A Cash Conversion System Perspective.

Benjamin Melamed, Rutgers Business School, USA.

A cash conversion system is a supply chain where funds are used to purchase inventory, and sales convert inventory back into funds. This work presents a new valuation method for cash conversion systems that simultaneously captures their endogenous product-level operating characteristics and cash flow dynamics using transient Markov Renewal processes. To address the computational challenges of Markov Renewal processes with a continuous state component, three discrete-state approximating processes are constructed. Specifically, these constitute an upper bound, a lower bound, and a convex combination thereof, which serve to approximate a cash conversion system's profitability and risk metrics. Numerical experiments confirm the efficacy and efficiency of this approximation methodology.

Keywords: approximations, business valuation, cash conversion system, Markov Renewal process, transient analysis.

Back to the time schedule.

Bandit Procedures for Designing Patient-Centric Clinical Trials.

Peter Jacko, Berry Consultants and Lancaster University, UK.

Multi-armed bandit problems (MABPs) are a special type of optimal control problem that has been studied in the fields of operations research, statistics, machine learning, economics, and others. It is a framework well suited to model resource allocation under uncertainty in a wide variety of contexts. The use of bandit models to optimally design clinical trials is one of the typical motivating application, for designing the so-called patient-centric trials, which would take into account the benefit of the in-trial patients. Nevertheless, the resulting theory has had little influence on the actual design of clinical trials. Contrary to similar learning problems arising for instance in digital marketing where interventions can be tested on millions of users at negligible cost, clinical trials are about "small data", as recruiting patients is remarkably expensive and (in many cases) ethically challenging. Due to the focus on small sizes, we do not resort to the use of the normal distribution to approximate a binomial distribution which is a common practice for large samples either "for simplicity" or "for ease of computation". We evaluate and compare the performance of a variety of operations research and machine learning procedures for the finite-horizon MABP, including the traditional and still dominant clinical trial design choice – equal fixed randomization – and interpret them in the context of designing clinical trials. Our results illustrate how bandit approaches could offer significant advantages, mainly in terms of allocating more patients to better interventions, but still pose important inferential challenges, particularly in terms of their resulting lower statistical power, potential for bias in estimation and existence of closed-form test distributions or asymptotic theory. We illustrate some promising modifications to bandit procedures to address power and bias issues, and we reflect upon the open challenges that remain for an increased uptake of bandit models in clinical trials. This is a joint work with Sofia S. Villar (MRC Biostatistics Unit, University of Cambridge).

Keywords: bandit problem, dynamic resource allocation, index rules, clinical trials.

Back to the time schedule.

Multilateral Decisions with Stopping Strategies.

Krzysztof Szajowski, Wroclaw University of Science and Technology, Poland.

The independent problem of the choice of the moment of action in a bilateral problem appeared in Dynkin (1969). The formulation of a stopping game assumes that there are two or more players who take turns making decisions, and each player has the option to either continue or stop the game at their turn. The players must make their decisions based on the information they have about the game's state and their opponent's behavior. The players may have different objectives, and the outcome of the game depends on the sequence of decisions made by the players. The assumption that makes the stopping game problem interesting is that the players have incomplete information about each other's preferences or strategies. This means, that each player must take into account: the possibility that their opponent will stop the game at any given turn, and the outcome of the game may depend on the player's ability to anticipate and react to their opponent's actions. The players may also face uncertainty about the future rewards or payoffs associated with continuing or stopping the game, which further complicates the decision-making process. Therefore, already in the formulation of the model, there is a need to take into account differences between players in terms of access to information or willingness to cooperate. Their individual preferences must be included in the definition of a rational conflict resolution. The contributions of professors Isaac Sonin and Michael Katehakis will be featured. Examples of applied solutions in this area are the subject of this report.

Keywords: Zero-sum game, no-zero-sum game, optional stopping, equilibrium, filtration.

Back to the time schedule.

Break

Session 3. 11.30 - 12.45 PM. Stella Kapodistria, Eindhoven University of Technology, Chair.

Boole's Probability Bounding Problem and the Zero-One Lemma.

Endre Boros, Rutgers University, USA.

We recall an old result of Renyi (1961) and show that numerous probability inequalities can be easily derived with its help. We also show some new results obtained with the help of this powerful and still very simple lemma. Joint work with Joonhee Lee (RUTCOR and MSIS, Rutgers University).

Keywords: Boole's problem, probability inequalities, bounding the union.

Back to the time schedule.

Optimal Stopping and Technical Analysis.

Saul Jacka, University of Warwick, UK.

Technical Analysis is a collection of investment policies based on the history of price processes. It is widely used by institutional investors despite conflict with the Efficient Markets Hypothesis. In this talk we'll discuss a very general model of a stock price which is designed to analyse the viability of a form of technical analysis known as the support and resistance line method.

Keywords: Optimal stopping; smooth pasting.

Back to the time schedule.

Integrated Learning and Decision Making Beyond Markov Settings.

Stella Kapodistria, Eindhoven University of Technology, The Netherlands.

In Markov decision processes, the archetypical point of departure is the Bellman optimality equations. For such problems, these equations rely on the underlying stochastic model being a Markov chain. However, the analysis and modelling of real data rarely produces a simple Markov chain model. Instead, it tends to result in history-dependent models (e.g., autoregressive models) or models with partial information. This necessitates two extensions of the current solution framework. Firstly, that the theoretical foundation of the Bellman optimality equations and the theory around them is extended to non-Markov settings. Secondly, that the learning of the underlying model is integrated with the decision analysis. Moreover, even within Markov settings, it is of paramount importance to create decision making approaches that on the one hand come with convergence guarantees and that on the other hand efficiently scale even for large computational problems without exhibiting the curse of dimensionality. In this talk, by focusing on a key Operations Research application, we will demonstrate that overcoming these challenges is tedious but not impossible.

Keywords: Operations Research; Markov Decision Process; Reinforcement Learning.

Back to the time schedule.

Session 4. 12.45 - 2.00 PM. Pavlo Kasyanov, National Technical University of Ukraine, Chair.

Reflections on Adaptation and Exploration.

Benjamin Van Roy, Stanford University, USA.

I will reflect on work in this area to which Professor Katehakis has greatly contributed.

Keywords: bandits, reinforcement learning.

Back to the time schedule.

Policy-Space Collapse in Two-Sided Bipartite Matching Models.

Rhonda Righter, UC Berkeley, IEOR, USA.

In two-sided matching models items arrive to one side and wait to be matched with items on the other side. Examples include buyers and sellers, patients and organs, and drivers and riders. Typically, there are multiple classes on each side with compatibility constraints that can be characterized by a bipartite compatibility graph. I will give conditions under which all position-based matching policies (e.g., first-come first-matched – among compatible matches, last-come first-matched, random matching) give the same transient distributions, and, for stable systems, the same stationary distributions. Examples in which policy-space collapse occurs are job and server assignment in redundancy-d models and idle server assignment for heterogeneous servers and exchangeable compatibilities.

Keywords: Bipartite matching models, parallel server systems, compatibility constraints, policy-space collapse.

Back to the time schedule.

Markov Decision Processes with Incomplete Information and Semiuniform Feller Transition Probabilities.

Pavlo Kasyanov, National Technical University of Ukraine, Ukraine.

This talk deals with control of partially observable discrete-time stochastic systems. It introduces and studies Markov Decision Processes with Incomplete Information and with semiuniform Feller transition probabilities. The important feature of these models is that their classic reduction to Completely Observable Markov Decision Processes with belief states preserves semiuniform Feller continuity of transition probabilities. Under mild assumptions on cost functions, optimal policies exist, optimality equations hold, and value iterations converge to optimal values for these models. In particular, for Partially Observable Markov Decision Processes the presented results imply new and generalize several known sufficient conditions on transition and observation probabilities for weak continuity of transition probabilities for Markov Decision Processes with belief states, the existence of optimal policies, validity of optimality equations defining optimal policies, and convergence of value iterations to optimal values.

Keywords: Markov decision process, incomplete information, semiuniform Feller transition probabilities, value iterations, optimality equations.

Back to the time schedule.

Break

Session 5. 2.30 - 4.10 PM. Csaba Szepesvari, University of Alberta, Chair.

A Unifying Reinforcement Learning Framework for Sequential Estimation, Bayesian Optimization, and Adaptive Control.

Dimitri Bertsekas, Arizona State University, USA.

We provide a unifying approximate dynamic programming framework that applies to a broad variety of problems involving sequential estimation. We consider first the construction of surrogate cost functions for the purposes of optimization, and we focus on the special case of Bayesian optimization, using the rollout algorithm and some of its variations. We then discuss the more general case of sequential estimation of a random vector using optimal measurement selection, and its application to problems of stochastic and adaptive control. We distinguish between adaptive control of deterministic and stochastic systems: the former are better suited for the use of rollout, while the latter are well suited for the use of rollout with certainty equivalence approximations. As an example of the deterministic case, we discuss sequential decoding problems, and a rollout algorithm for the approximate solution of the Wordle and Mastermind puzzles, recently developed in a joint paper with S. Bhambri and A. Bhattacharjee.

Keywords: Bayesian Optimization, Adaptive Control.

Back to the time schedule.

Black-box Quantile Optimization via Finite Difference-based Gradient Approximation.

Jiaqiao Hu, Stony Brook University, USA.

We consider quantile optimization problems under a general black-box setting. We propose two new iterative multi-time scale stochastic approximation type of algorithms. The first algorithm uses an appropriately modified finite-difference-based gradient estimator that requires 2d+1 samples of the black-box function per iteration of the algorithm, where d is the number of decision variables. The second algorithm employs a simultaneous-perturbation-based gradient estimator that uses only three samples for each iteration regardless of problem dimension. We show the almost sure convergence of both algorithms and establish their rates of convergence. Numerical results are also reported to illustrate and compare the performance of the algorithms with alternative methods.

Keywords: quantile optimization, stochastic approximation.

Back to the time schedule.

Optimal Allocation of Resources for Network Systems.

Sergei Schreider, Rutgers Business School, USA.

In my presentation I will demonstrate several cases of network optimization modelling for optimal allocation of resources. The major examples of such applications will be related to the water allocation and optimal gas supply in Eastern Australia.

Keywords: Optimization, network system, resource management.

Back to the time schedule.

The Curse of Dimensionality and Function Approximation in Reinforcement Learning.

Csaba Szepesvari, University of Alberta, Canada.

Reinforcement learning (RL) has been an active field of research for over 60 years, with its roots in the work of Richard Bellman and his colleagues. Their pioneering work introduced dynamic programming and function approximation as a solution to the curse of dimensionality, a problem that has plagued RL since its inception. Since then, the field has seen significant progress from various research communities. In this talk, we will explore recent advances in overcoming the curse of dimensionality using function approximation. We will discuss the latest techniques and approaches that have been developed to address this challenge, as well as the fundamental limits of this approach. We will also highlight the key challenges that still need to be addressed in the field of RL.

Keywords: reinforcement learning; function approximation.

Back to the time schedule.

Break

Session 6. 4.40 - 6.20 PM. Yao Zhao, Rutgers University, Chair.

P-values, heterogeneity, and bias for rating risk factors: insights from umbrella reviews of observational studies.

John P.A. Ioannidis, Departments of Medicine, of epidemiology and Population Health, of Biomedical Data Science, and of Statistics and Meta-Research Innovation Center at Stanford (METRICS), USA.

Evidence on risk factors in observational studies is fragmented and notoriously debated as to its validity and causal inference potential. Most published analyses choose and report candidate risk factors based on p-value thresholds and the same applies to meta-analyses of multiple studies thereof. A growing literature suggests that P-values, even when very low, have little ability of identifying reliably important risk factors, let alone support causality. Empirical evidence collected from many thousands of systematic reviews and meta-analyses may nevertheless offer some empirical calibration of different criteria based on P-values, in combination and in juxtaposition to criteria that evaluate consistency and heterogeneity as well as common tests of bias in the published literature, such as excess significance tests and small study effect asymmetry tests.

Keywords: P-values, observational studies, risk factors, bias.

Back to the time schedule.

COSIMLA: Efficiently COmbining SIMulation with Linear Algebra.

Peter Glynn, Stanford University, USA.

For high accuracy computations involving discrete state space Markov chains in discrete and continuous time, using numerical linear algebra is the preferred approach when the state space is of tractable size. When the state space is large, the linear system must be truncated in some way, thereby incurring error that is typically of an unknown magnitude. In this talk, we provide the first algorithm that solves a linear system on the "core" of the state space, and uses simulation (sparingly) to compute the contributions to the probability or expectation that come from excursions outside the "core" set, thereby providing high accuracy solutions that are computationally efficient.

Keywords: Markov chains, numerical methods, simulation.

Back to the time schedule.

Supply Chain Analytics – From Problem Solving to Problem Discovery.

Yao Zhao, Rutgers University, USA.

I will discuss how to use data analytics to identify and discover problems in supply chains, such as, inventory analytics, sourcing analytics and competitive intelligence. The focus is on transforming data into insights - you may find that data-driven problem discovery is not only highly valuable, but also sophisticated and revealing.

Keywords: Data analytics, problem discovery, supply chains.

Back to the time schedule.

Finding a Population Having the Better Distribution.

Sheldon Ross, University of Southern California, USA.

Suppose there are an infinite number of populations, with each population having an associated probability distribution, and that sampling from a population incurs a cost 1 and yields the value of a random variable having the population distribution. Suppose there are only 2 possible population distributions, with densities f and g. Suppose that f is the density of population i with probability p(i) when i=1,...,k and with probability p when i>k. At each stage one selects one of

the populations to sample from, with this continuing until there is a population whose posterior probability of having density f is at least some specified value. The objective is to find a sampling strategy that minimizes the expected cost until this occurs. We also consider the problem where one may stop at any time and choose a population, with a reward of R earned if that population has distribution f.
Keywords: optimal stopping, sequential sampling, Gittins index.
Back to the time schedule.

Break

Open Session. 6.45 - 7.45 PM. Sheldon Ross, University of Southern California, Facilitator.
Back to the time schedule.

Saturday April 22, 2023.
Back to the time schedule.

Session 1. 8.30 - 9.45 AM. Apostolos Burnetas, National and Kapodistrian University of Athens, Chair.

SIEVE: A Space-Efficient Algorithm for Viterbi Decoding.
Panagiotis Karras, Aarhus University, Denmark.
Can we get speech recognition tools to work on limited-memory devices? The Viterbi algorithm is a classic dynamic programming (DP) solution used to find the most likely sequence of hidden states in a Hidden Markov Model (HMM). While the algorithm finds universal application ranging from communication systems to speech recognition to bioinformatics, its scalability has been scarcely addressed, stranding it to a space complexity that grows with the number of observations. In this paper, we propose SIEVE (Space Efficient Viterbi), a reformulation of the Viterbi algorithm that eliminates its space-complexity dependence on the number of observations to be explained. SIEVE discards and recomputes parts of the DP solution for the sake of space efficiency, in divide-and-conquer fashion, without incurring a time-complexity overhead. Our thorough experimental evaluation shows that SIEVE is highly effective in reducing the memory usage compared to the classic Viterbi algorithm, while avoiding the runtime overhead of a naïve space-efficient solution.
Keywords: HMMs, Viterbi decoding.
Back to the time schedule.

Recursive Computation of Equilibrium and Optimal Strategies in Unobservable Feed-forward Queueing Networks with Delay-Sensitive Customers.
Apostolos Burnetas, National and Kapodistrian University of Athens, Greece.
In this talk we consider the problem of determining equilibrium traffic intensities, and thus congestion and delays, in service networks with forward flows. Specifically, we consider a feed-forward open Jackson network with given external arrival rates in each queue. Customers are strategic and make join/balk as well as routing decisions based on a service reward/delay cost tradeoff. External customers arriving to a queue decide whether to join or balk. Customers who finish service in a queue decide to join one of the immediately following queues or balk. The number of servers, the exponential service rate, the service reward and the delay cost rate are generally queue-dependent. We consider the unobservable case, where customers do not have any information on the system state, however they know the network structure as well as all system parameters. The customer's objective is to maximize the expected total net benefit from the entire route she will follow in the system. We consider customer equilibrium strategies for the unobservable case. We develop recursive equations for the customer expected benefit function and show that under fairly general conditions there exist unique, generally mixed, equilibrium strategies. In the special case of a two-stage network with two queues in the first stage leading to a single queue in the second stage, we analyze the form of equilibrium strategies under two cases: (i) a customer selects which, if any, queue to join in the first stage and is then required to be served in the second stage queue as well and, (ii) a customer first decides which queue, if any, to join in the first stage, and then whether to proceed to the second stage or abandon the system. We explore numerically the effect of the abandonment option on the equilibrium system loads and delays. We next consider the problem of optimally admitting/routing customers within the network so as to maximize the expected total benefit of all customers. In order to compare equilibrium and optimal strategies, we also consider the unobservable, i.e., state-independent case of the social optimization problem. We

present analogous recursive expressions that lead to a dynamic programming formulation. We also discuss the relationship between equilibrium and optimal strategies.
Keywords: Strategic Customers; Queueing Networks; Equilibrium Strategies; Social Optimal Strategies.
Back to the time schedule.

Preferences, Risk-Neutrality, and Risk-Sensitive MDPs.

Matthew Sobel, Case Western Reserve University, USA.

A binary preference relation on a real vector space satisfying four (natural) axioms is shown to induce a utility function composed of a linear function to the reals and a weakly monotonic function. The key axiom is decomposition, and the utility function can be taken to be linear if and only if this axiom's converse is also satisfied. Important consequences follow for decision trees, risk-sensitive discounted Markov decision processes, and the discounted utility model in economics. Since the four axioms imply that preferences correspond to discounting, the four axioms without the converse imply that preferences are consistent with discounting without risk neutrality. (Joint research with James C. Alexander).
Keywords: Preference, Utility, Discounting, MDP.
Back to the time schedule.

Session 2. 9.45 - 11.00 AM. Eugene Feinberg, Stony Brook University, Chair.

Predictive Analytical Models for Data Sciences. Stochastic Intensity Function and Stochastic Monitoring Indicator of the Power Law Process.

Chris Tsokos, University of South Florida, USA.

Developing real data driven predictive analytical models that identify the risk factors and interactions that cause a phenomenon of interest in life sciences, global warming, business and economics. Two entities of the Power Law Process that monitor the behavior of the response of a non stationary signal.
Keywords: Power Law Process, Stochastic Intensity, Monitoring Indicator.
Back to the time schedule.

Sequential Optimization of CVaR.

Eugene Feinberg, Stony Brook University, USA.

This talk describes theoretical foundations and algorithms for optimizing Conditional Value at Risk (CVaR) for a discounted total-cost Markov Decision Process (MDP) with finite state and action sets. This CVaR optimization problem can be reformulated as a Robust MDP(RMDP) with compact state spaces. States in these RMDPs are the original states of the problems augmented with tail risk levels, and the Decision Maker (DM) knows only the tail risk level at the initial state and time. Thus, in order to find an optimal policy following this approach, the DM needs to solve an RMDP with incomplete state observations because after the first move the DM observes the states of the system, but the tail risk levels are unknown. We show that for the CVaR optimization problem the corresponding RMDP can be solved by using the methods of convex analysis. The talk presents an algorithm for computing and implementing an optimal CVaR policy by using the value function for the version of this RMDP with completely observable tail risk levels at all states. This algorithm and the major results of the talk cover a more general problem of optimizing sum of a mean and CVaR for possibly different cost functions. This talk is based on a joint paper with Rui Ding.
Keywords: Conditional Value at Risk, Robust Markov Decision Process, optimality equation, computing an optimal policy.
Back to the time schedule.

The potential of MCMC, with Quantum computing.

Tze Lai, Stanford University, USA.

TBA
Keywords: TBA
Back to the time schedule.

Break

Session 3. 11.30 - 12.45 PM. Farid Alizadeh, Rutgers University, Chair.

Risk-Averse Control of Markov Systems with Value Function Learning.

Andrzej Ruszczynski, Rutgers University, USA.

We consider a control problem for a finite-state Markov system whose performance is evaluated by a coherent Markov risk measure. For each policy, the risk of a state is approximated by a function of its features, thus leading to a lower-dimensional policy evaluation problem, which involves non-differentiable stochastic operators. We introduce mini-batch transition risk mappings, which are particularly suited to our approach. We propose a robust learning algorithm for Markov policy evaluation. Finally, we discuss structured policy improvement in the feature-based risk-averse setting. The considerations are illustrated with a supply chain problem and an underwater robot navigation problem.

Keywords: Dynamic Risk Measures, Reinforcement Learning, Function Approximation.

Back to the time schedule.

Strategies to Outsource Acquisition and Retention: the Impact of Effort Comparison and Risk Preference.

Steve Bin Zhou and Xinxin Hu, University of Houston Downtown, USA.

This paper analyzes how the brand company chooses the optimal outsourcing strategy of customer acquisition and retention under the service supplier's risk preference and effort comparison effect. We examine three outsourcing strategies: outsourcing customer acquisition (AO), outsourcing customer retention (RO), and outsourcing both customer acquisition and retention (TO). We evaluate the supplier's incentives with different risk preferences, i.e. risk neutral and risk averse, under the conditions that consumers could compare marketing efforts of customer acquisition and customer retention, but they also have a probability not to respond to any effort. Our main result reveals that the incentive mechanism can help the brand company obtain higher profits under a given outsourcing strategy. In addition, the service supplier's risk preference affects the brand company's choice of service outsourcing strategy. If the supplier is risk neutral, the brand company can outsource both customer acquisition and retention. If the supplier is risk averse, the company would be better off outsource only customer acquisition.

Keywords: service supply chain; OR in marketing; incentive mechanism; risk preference.

Back to the time schedule.

Learning with Ambiguous Data, Baked-in Regularization in Regression.

Farid Alizadeh, Rutgers University, USA.

In many applications, the input data come with errors, and the numbers and categories given are inaccurate. In such situations, we propose that the data be presented as probability distributions that depend on the observation mechanism. For instance, suppose in some data sets we want to store patients' weight and blood pressure. However, these features may change during the day. Furthermore, the scale on which weight is measured or the blood pressure gauge on which pressure is measured might be imprecise. So, instead of recording the patient's weight as a precise number x, we may alternatively say that the observed weight follows a distribution with a pdf, say the normal distribution with a given mean and variance; we call this the observational distribution of the datum. In general, a data set of N observations, with d input and m target variables, is represented by a set of N joint observational distributions. We examine the effects of this form of data, mainly when applied to linear and logistic regression methods. In particular, we show that introducing ambiguity to data will induce regularization without any need to determine hyperparameters.

Keywords: ambiguous data, regression, regularization.

Back to the time schedule.

Session 4. 12.45 - 2.00 PM. Jian Yang, Rutgers University, Chair.

Examining the Relationship between Managerial Ability and Firm Corporate Social Responsibility Performance in the Energy Industry: A Stochastic Data Envelopment and Panel Data Analysis.

Steve Bin Zhou and Ray Cao, University of Houston Downtown, USA.

In this paper, we utilized a two-step approach for exploring the relationship between managerial ability and firm sustainability (i.e., ESG) performance in the energy industry in US. We employed a stochastic DEA approach to derive the managerial efficiency for the managerial ability. Then, we tested the hypotheses via the fixed-effects panel data analysis to examine the associations between the managerial ability and different aspects of the ESG performance. Our main findings suggest that

energy firms in US need to either hire top executives with excellent managerial abilities or enhance their top executives' managerial abilities for these firms to enhance ESG performances.
Keywords: Managerial Ability; Corporate Social Responsibility; Stochastic DEA; Panel Data Analysis.
Back to the time schedule.

Partition-based Stability of Coalitional Games.

Jian Yang, Rutgers University, USA.

We are concerned with the stability of partition-allocation pairs in a coalitional game, i.e., a transferable-utility (TU) cooperative game. First, the concept of core can be weakened so that the blocking of changes is limited to only those with multilateral backings. This principle of consensual blocking can then be applied to partition-allocation pairs. Each such pair is made up of a partition of the grand coalition and a corresponding allocation vector whose coalition-wise components are efficient for the various constituent coalitions of the given partition. In a roundabout way, the resulting stability concept is shown to be universal, meaning that any game, no matter how ``poor'' it is, has its fair share of stable solutions. Via bottom-up constructions, we can reach two tighter and yet still provably-universal stability notions, one local and another global. Local stability can even be reached via a greedy multi-step process.
Keywords: Cooperative Game; Coalition; Partition; Stability.
Back to the time schedule.

Privacy-preserving Data Sharing and Analytics.

Jaideep Vaidya, Rutgers University, USA.

In the current digital age, data is continually being collected by organizations and governments alike. While the goal is to use this data to derive insight and improve services, the ubiquitous collection and analysis of data creates a threat to privacy. In this talk, we examine the underlying challenges in enabling data sharing and analytics in a privacy-preserving manner. This requires a combination of technological advances to ensure both process and output privacy as well as socio-cognitive approaches to ensure the widespread use and deployment of such work. We discuss some of our recent research results on identifying anomalies and crowdsensing data in this context.
Keywords: Privacy, Data Sharing, Analytics.
Back to the time schedule.

Break

Session 5. 2.30 - 4.10 PM. Jim (Junmin) Shi, New Jersey Institute of Technology (NJIT), Chair.

Variance-adaptive Regret Bounds for Control of Ergodic MDPs.

Sadegh Talebi, University of Copenhagen, Denmark.

In this talk, we consider adaptive control of finite Markov decision processes (MDPs) under the average-reward criterion, in the regret minimization setting, where the performance of the adaptive control strategy in question is measured through the notion of regret. We consider ergodic MDPs, where the transition law induced by any stationary policy is irreducible. An asymptotic problem-dependent lower bound as well as an algorithm asymptotically achieving the lower bound were presented in the seminal work by Burnetas & Katehakis (1997). These bounds grow logarithmically with time, where the leading constant is a complex MDP-dependent constant. In this talk, we present a simplification of the leading constant in the regret lower (and upper) bound, which leads to making appear more explicit MDP-dependent quantities. One key quantity here is the variance (with respect to next-state distributions) of the optimal bias function appearing in the Bellman optimality equation. Further, we discuss the KL-UCRL algorithm and present a non-asymptotic regret bound (scaling as the root of time horizon) for it that depends on the same variance terms.
Keywords: adaptive control; Markov decision process; regret; Bellman optimality equation; transportation inequalities.
Back to the time schedule.

Robust, Scalable and Explainable Analytics for Biomedical Applications.

Dimitris Metaxas, Rutgers University, USA.

Over the last 30 years, we have been developing a general, scalable, computational learning and AI framework that combines principles of computational learning, neural nets, sparse methods, mixed norms, dictionary learning, and deformable modeling methods. This framework has been used for resolution of complex large scale problems in biomedical image analysis and collaborations with the pharmaceutical industry. Our learning methods allow the discovery of complex features, shapes, relationships, disease diagnosis for many types of clinical and preclinical applications. We will present segmentation, registration, tracking and disease recognition methods and their applications to cardiac analytics, cancer diagnosis, body fat

estimation and cell tracking. Finally, we will show novel AI methods that offer explainability in machine learning to provide further insights into learning-based decision making and diagnosis.
Keywords: AI, Machine Learning, Medical Domain Knowledge, Biomedical Data.
Back to the time schedule.

Blockchain Technology Optimal Adoption for Matching Random Supply with Uncertain Demand.

Jim (Junmin) Shi, New Jersey Institute of Technology (NJIT), USA.

Blockchain has been embraced as a disruptive technology, but very little research examines at its impact and adoption issues from the Supply Chain optimization perspective. Motivated by the prevailing challenges faced by AgriLedger platform for Haitian Mango supply chain, this paper investigates the impact of blockchain technology (BCT) on Supply Chain Management (SCM) and then focuses on the fundamentally inherent adoption issues. In particular, this study aims to first examine the impact of BCT adoption in the presence of random supply and uncertain demand, and further develop the optimal adoption solution. We consider a stochastically dynamic programming model, where a firm seeks to maximize the total expected discounted profit, by jointly managing (i) blockchain adoption, (ii) production and ordering decisions, and (iii) dynamic pricing and selling. We first show that the deployment of BCT can assist firms in reducing order quantities, lowering selling prices and reducing target-inventory levels. It is also shown that volatility of either supply or demand lowers the expected profit. The analysis is robust with some major extensions, such as lost-sales of demand and random capacity. Our numerical study accumulates useful managerial insights. For example, subject to tech-savvy customer behavior, some types of goods (e.g., credence goods and experience goods) greatly benefit from the adoption of BCT, but it may not prove beneficial to leverage BCT for certain others (e.g., search goods). Considering the lifecycle of a typical good, it is recommended to adopt BCT as early as possible and to adopt it to a higher degree at an earlier stage.
Keywords: Blockchain Technology, Supply Chain, Random Yields, Stochastic Monotonicity, AgriTech.
Back to the time schedule.

Comparing Centralized and Decentralized Matching in Labor Markets.

Karti Puranam, Rutgers University, USA.

In this paper the problem of efficiently matching applicants to jobs is studied. Two mechanisms for matching which are dubbed centralized and decentralized mechanisms are compared. The payoff to the Labor Market Intermediary (LMI) which facilitates the matching in both mechanisms is studied. A discussion of when each of the two mechanisms are optimal is provided. Numerical simulations and their results are also presented.
Keywords: Market Design, Two-sided matching, labor market intermediary.
Back to the time schedule.

Break

Session 6. 4.40 - 6.20 PM. Javier Cabrera, Rutgers University, Chair and Organizer.

How Much Does It Cost?

Victor de la Pena, Columbia University, USA.

How much does it cost a decisionmaker to base her payoff on interdependent, biased information sources? This question is relevant in economics, statistics, and politics. When there are many information sources, their dependence may be unknown or uncertain, which creates multivariate ambiguity. One approach to answer our leading question involves decoupling inequalities. We present new inequalities which hold for any type of dependence as well as some applications.
Keywords: decoupling inequalities, cost.
Back to the time schedule.

Estimate excess COVID deaths in the US by cause of death.

Davit Sargsyan, Rutgers University and Johnson & Johnson Pharmaceutical, USA.

Causes of death incidence in America fluctuates year by year. Only deaths with an autopsy are sure to be correct, the other deaths are classified by the most probable cause without any test to prove it. Besides, healthcare conditions, economics, natural disasters, and drug overdose creates fluctuations on the American death landscape. COVID-19 has brought further uncertainties given the additional COVID test necessary to achieve the correct cause of death of the deceased Americans. The COVID deaths data don't catch everyone whose life was shortened by the pandemic and add other people whose primary reason for dying was not COVID. This paper shows how to detect excess deaths using causes of death data. The analysis is

conducted at the cause of death level to identify the over and under reported deaths by COVID or non COVID. Joint work with Davit Sargsyan, Johnson & Johnson.
Keywords: Covid deaths, predictive model.
Back to the time schedule.

Two Stage Errors and Variables with Clinical Applications.

Yajie Duan, Rutgers University, USA.

Risk assessment models are widely used in clinical settings, such as CHADS2[1] for predicting one-year risk of stroke after the initial atrial fibrillation diagnosis, and HAS-BLED[2] for bleeding risk assessment in patients taking blood thinners. However, there are few or no risk assessment models focused on the relationship between multiple risks. For instance, patients with high risk of stroke are usually treated with anticoagulants, but anticoagulants inherently increase the risk of bleeding. This suggests that risks of stroke and bleeding should be modeled jointly to decide whether or not to treat with anticoagulants. To help with this decision making, a novel framework with a two-stage Deming regression model[3] is proposed to predict stroke and bleeding risks simultaneously and provide the relationship between them. Deming regression is an errors-in-variables model which tries to find the line of best fit for variables considering errors on both the x- and the y- axis.[4] The first stage of the proposed model estimates the risks of stroke and bleeding with their corresponding standard errors, and after suitable transformations for linearizing the relationship, in the second stage, a Deming regression line is fitted to the transformed risks. The model accounts for the prediction errors with known variance and measurement errors, which is much more reasonable than using traditional regression models directly. The proposed method could be applied to analysis between multiple risks, especially those weighed against each other such as risks of illness and side effects of drugs. In addition, patients' subjective opinions about the outcomes will also be included in this innovated multiple risk assessment system. Joint work with Javier Cabrera.
Keywords: Deming regression, risk assessment, stroke, bleeding.
Back to the time schedule.

Weighted Probabilistic Distance Clustering for Big Data.

Debopriya Ghosh, Johnson & Johnson Pharmaceutical, USA.

Big Data introduces statistical and computational challenges. Traditional algorithms do not scale to massive datasets. Reducing large sample size can greatly improve performance of algorithms. In this talk, I will present a new probabilistic, iterative method for clustering weighted data, using soft assignments of points to clusters with membership probabilities depending on distances and cluster sizes. We refer to it as weighted probabilistic distance (w-PDQ) clustering, where Q stands for cluster size. The novel aspect of the proposed method is the ability to handle weighted reduced data, which makes it suitable for clustering large datasets. Experiments on simulated and real data demonstrate that the weighted probabilistic distance clustering approach performs favorably to other model-based clustering approaches. In addition, the approach is robust to outliers and computationally efficient as it does not require computing complex density functions.
Keywords: Big Data, clustering, weighted clustering.
Back to the time schedule.

Break

Open Session. 6.45 - 7.45 PM. Eugene Feinberg, Stony Brook University, Facilitator.

Back to the time schedule.

The conference is sponsored by Leiden University, and the program committee is: Odysseas Kanavetas (chair, Leiden University), Eugene Feinberg (Stony Brook University), Sheldon Ross (University of Southern California) and Floske Spieksma (Leiden University).