Uncertainty Management and Machine Learning in Engineering Applications
November 16, 2020
Speakers

10:00AM – 10:30AM – Pawel Polak, Stony Brook University (pawel.polak@stonybrook.edu)
Stan Uryasev, Stony Brook University (stanislav.Uryasev@stonybrook.edu)
Teng Chen, Stony Brook University (teng.chen@stonybrook.edu)

Title: Classification and Severity Progression Measure of COVID-19 Patients Using Proteomic and Metabolomic Sera

Abstract: Early detection and effective treatment of severe COVID-19 patients remain one of the major challenges during the current pandemic. Analysis of molecular changes in blood samples of severe patients is one of the promising approaches to this problem. From 75 most relevant proteomic and metabolomic biomarkers selected by Shen et al. (2020), we identify several *pairs of biomarkers* (some of them after additional nonlinear spline transformation) that are highly effective in classifying severe COVID-19 cases. The performance of these pairs is evaluated in-sample, in a cross-validation exercise, and in an out-of-sample analysis on an independent dataset. Our findings can help medical experts to identify small groups of biomarkers that can be used to construct a cost-effective, short-term test for patients screening and a measure of severity progression. This is a joint work with Stan Uryasev and Teng Chen.

10:30AM – 11:00AM – Craig Vineyard, Sandia National Laboratories (cmviney@sandia.gov)

Title: Neural Network Approaches for Enabling Automatic Target Recognition

Abstract: Automatic target recognition (ATR) is an exploitation algorithm for detection via a remote sensor. This high consequence decision making process must deal with sensor induced challenges, difficult to detect targets, and limited training data. With deep neural networks (DNNs) offering many state-of-the-art solutions to computer vision tasks, neural networks are once again being revisited for ATR processing. In this talk, we characterize and explore a suite of neural network architectural topologies for enabling ATR. In doing so, we assess how different architectural approaches impact performance as well as consider the associated computational costs. And lastly, we introduce some insights into further understanding their operation as a high consequence application.

11:00AM – 11:30AM – Warren Dixon, University of Florida (wdixon@ufl.edu)

Title: Multiple Timescale Deep Learning

Abstract: A Deep Neural Network (DNN) adaptive control architecture is presented for general uncertain nonlinear dynamical systems to track a desired time-varying trajectory. A Lyapunov-based method is leveraged to develop adaptation laws for the output-layer connections of a DNN model in real-time while a data-driven supervised learning algorithm is used to update the inner-layer connections of the DNN. Specifically, the output-layer connections of the DNN are estimated using an unsupervised learning algorithm to provide responsiveness and guaranteed tracking performance with real-time feedback. The inner-layer connections of the DNN are trained with collected data sets to increase performance, and the adaptation laws are updated once a sufficient amount of data is collected. A nonsmooth Lyapunov-based analysis is used to prove semi-global asymptotic tracking of the desired trajectory. A numerical simulation example is included to validate the results.

11:30AM – 12:00PM – Oleg Prokopyev, University of Pittsburgh (droleg@pitt.edu)

Title: A Mixed-Integer Fractional Optimization Approach to Best Subset Selection

Abstract: We consider the best subset selection problem in linear regression, i.e., finding a parsimonious subset of the regression variables that provides the best fit to the data according to some predefined criterion. We are primarily concerned with alternatives to cross-validation methods that do not require data partitioning and involve a range of information criteria extensively studied in the statistical literature. We show that the problem of interests can be modeled using fractional mixed-integer optimization, which can be tackled by leveraging recent advances in modern optimization solvers. The proposed algorithms involve solving a sequence of mixed-integer quadratic optimization problems (or their convexifications) and can be implemented with off-the-shelf solvers. We report encouraging results in our computational experiments, with respect to both the optimization and statistical performance. This is a joint work with Andres Gomez (University of Southern California).

12:00PM – 12:30PM – Break

12:30PM – 1:00PM – Andrzej Ruszczynski, Rutgers University (rusz@business.rutgers.edu)

Title: A Single Time-Scale Stochastic Approximation Method for Nonsmooth and Nonconvex

Stochastic Composition Optimization

Abstract: We study nested stochastic optimization problems in which the objective function is a composition of several non-smooth and non-convex functions whose exact values and generalized derivatives are not available and only their stochastic estimates can be obtained. We illustrate the relevance of this problem by applications to stochastic variational inequalities, reinforcement learning, and risk-averse optimization.

As a tool of our analysis, we prove a chain rule on a path for functions which are differentiable in a generalized way. Then we analyze a stochastic optimization algorithm with direction averaging for such functions and illustrate its operation on a ReLU neural network. For the composition problem, we prove convergence in the general nonsmooth case and discuss the rate of convergence in the smooth case. Finally, we illustrate the efficacy of the method in distributionally robust learning.

1:00PM – 1:30PM – Drew Kouri, Sandia National Laboratories (dpkouri@sandia.gov)

Title: Design of Experiments for Superquantile Regression

Abstract: Traditional least-squares regression is often insufficient for modeling high-consequence applications because it tracks the mean behavior of the system. In this talk, we investigate the use of superquantile regression to instead track the average value-of-risk. Unfortunately, constructing accurate statistical models of critical system responses often requires an enormous amount of experimental data, which can be expensive and time consuming. To address this issue, we formulate an optimal experimental design problem that determines the "best" allocation of experiments by minimizing some measure of estimation or prediction variance. To this end, we first analyze the large-sample statistics of superquantile regression, demonstrating that the sequence of estimators is consistent and asymptotically normal. Using the associated asymptotic covariance matrix, we then formulate the optimal design problem. We solve the design problem using a primal-dual algorithm that was originally introduced for risk-averse optimization. We demonstrate our approach by designing experiments for direct field acoustic testing: a technique used to test engineered structures in vibration environments by subjecting them to intense acoustic pressure.

1:30PM – 2:00PM – David Stracuzzi, Sandia National Laboratories (djstrac@sandia.gov)

Title: Models of Models: Recognizing and Managing the Uncertainties of Machine Learning in Engineering Applications

Abstract: Machine learning is a powerful approximation tool that is steadily becoming the method of choice for many engineering applications. However, customizing a learned model to the needs of a particular application requires navigating a complex space of design decisions about the data, modeling, and training approach. In this talk, we first examine this design space and then discuss methods for navigating through it. Along the way, we identify a few key insights and guidelines intended to minimize the risk of machine learning misuse.

2:00PM – 2:30PM – Eric Cyr, Sandia National Laboratories (eccyr@sandia.gov)

Title: A Layer-Parallel Approach for Training Deep Neural Networks

Abstract: Deep neural networks are a powerful machine learning tool with the capacity to "learn" complex nonlinear relationships described by large data sets. Despite their success training these models remains a challenging and computationally intensive undertaking. In this talk we will present a new layer-parallel training algorithm that exploits a multigrid scheme to accelerate both forward and backward propagation. Introducing a parallel decomposition between layers requires inexact propagation of the neural network. The multigrid method used in this approach stiches these subdomains together with sufficient accuracy to ensure rapid convergence. We demonstrate an order of magnitude wall-clock time speedup over the serial approach, opening a new avenue for parallelism that is complementary to existing approaches. Results for this talk can be found in [1,2].

[1] S. Guenther, L. Ruthotto, J. B. Schroder, E. C. Cyr, N. R. Gauger, Layer-Parallel Training of Deep Residual Neural Networks, SIMODs, Vol. 2 (1), 2020.

[2] E. C. Cyr, S. Guenther, J. B. Schroder, Multilevel Initialization for Layer-Parallel Deep Neural Network Training, arXiv preprint arXiv:1912.08974, 2019

2:30PM – 3:00PM – Break

3:00PM – 3:30PM – Suvrajeet Sen, University of Southern California (suvrajes@usc.edu)

Title: Stochastic Decomposition - Computational Advances over Three Decades

Abstract: Stochastic Decomposition (SD) is a computational approach to solving large scale Stochastic Linear (SLP) and Quadratic Programming (SQLP/SQQP) models. Starting with its inception, almost three decades ago, it has continued to set new computational benchmarks, as well as newer highly competitive convergence rates. Finally, we will present introduce the notion of compromise of decisions which provide reduced variance decisions This is joint work with many former students, including Harsha Gangammanavar (2014), Yifan Liu (2016) and Junyi Liu (2019)

3:30PM – 4:00PM – Bart van Bloemen Waanders, Sandia National Laboratories (bartv@sandia.gov)

Title: Hyper-Differential Sensitivity Analysis: Managing High Dimensional Uncertainty in Large-Scale Optimization

Abstract: Large-scale optimization is ubiquitous in scientific and engineering applications. The end goal in most applications is the solution of a design, control, or inverse problem, constrained by complex high-fidelity models and robust in the face of uncertainties. Achieving this goal is challenging for many reasons, most notably, the computational complexity of the models and their numerous sources of uncertainty. This talk introduces hyper-differential sensitivity analysis (HDSA) as a tool to assess the sensitivity of optimal solutions with respect to uncertainties. Differing from traditional sensitivity analysis, HDSA provides unique insights in the context of the optimization problem rather than simply the model. The mathematical foundations and scalable algorithms are presented and insightful examples demonstrate the utility of the proposed approach for nonlinear and multi-physics problems.

4:00PM – 4:30PM – Thomas Surowiec, Philipps University of Marburg, Germany (surowiec@mathematik.uni-marburg.de)

Title: Stability Analysis for a Class of Risk-Neutral PDE-Constrained Optimization Problems

Abstract: Stability analysis in stochastic optimization has long been used to investigate the behavior of both the optimal value and optimal solutions with respect to perturbations in the underlying measure. This is particularly important information for numerical approaches in which the probability measure is replaced using Monte-Carlo methods or by a data-based empirical measure. In such instances, it is crucial to understand such dependencies as we increase the sample size or pass to the large data limit. Due to the weak convergence properties of probability measures, there are no guarantees that the optimal values or solutions will exhibit any favorable properties such as continuity or convergence to a solution of a limiting optimization problem.

Taking as our inspiration recent work on PDE-constrained optimization under uncertainty as well as several well-known models from functional data analysis, we develop a stability analysis for a class of infinite-dimensional stochastic optimization problems using the method of probability metrics. We show that the optimal values are Lipschitz continuous with respect to a minimal information metric and subsequently, under further assumptions, with respect to certain Fortet-Mourier and Wasserstein metrics. The optimal solutions are shown to be at best Hölder continuous. The general results are then applied to a class of risk-neutral PDE-constrained optimization problems, which lead to an a priori error estimate. Finally, we provide several numerical results both to illustrate the theory and to point to potential limitations of the method of probability metrics.

4:30PM – 5:00PM – Alexander Shapiro, Georgia Institute of Technology (ashapiro@isye.gatech.edu)

Title: Computational Approaches to Solving Multistage Stochastic Programs

Abstract: In this talk we discuss computational approaches to solving convex stochastic programming problems. In some applications the considered multistage stochastic programs have a periodical behavior. We demonstrate that in such cases it is possible to drastically reduce the number of stages by introducing a periodical analog of the so-called Bellman equations, used in Markov Decision Processes and Stochastic Optimal Control. Furthermore, we describe a primal-dual variant of the Stochastic Dual Dynamic Programming algorithm, applied to the constructed periodical Bellman equations, and show numerical experiments for the Brazilian interconnected power system problem.

5:00PM – 5:30PM – Break

5:30PM – 6:00PM – Lars Ruthotto, Emory University (lruthotto@emory.edu)

Title: Machine Learning for High-Dimensional Optimal Transport

Abstract: In recent years, OT and ML have become increasingly intertwined. This talk presents new avenues for solving high-dimensional optimal transport (OT)problems using machine learning (ML).

The first part of the talk shows how neural networks can be used to efficiently approximate the optimal transport map between two densities in high dimensions. To avoid the curse-of-dimensionality, we combine Lagrangian and Eulerian viewpoints and employ neural networks to solve the underlying Hamilton-Jacobi-Bellman equation. Our approach avoids any space discretization and can be implemented in existing machine learning frameworks. We present numerical results for OT in up to 100 dimensions and validate our solver in a two-dimensional setting.

The second part of the talk shows how optimal transport theory can improve the efficiency of training generative models and density estimators, which are critical in machine learning. We consider continuous normalizing flows (CNF) that have emerged as one of the most promising approaches for variational inference in the ML community. Our numerical implementation is a discretize-optimize method whose forward problem relies on manually derived gradients and Laplacian of the neural network and uses automatic differentiation in the optimization. In common benchmark challenges, our method outperforms state-of-the-art CNF approaches by 1-2orders of magnitude during training and inference.

6:00PM – 6:30PM – Darinka Dentcheva, Stevens Institute of Technology (darinka.dentcheva@stevens.edu)

Title: Bias Reduction in Sample-Based Optimization

Abstract: We consider stochastic optimization problems which use observed data to estimate essential characteristics of the random quantities involved. Sample average approximation (SAA) or empirical (plug-in) estimation is one of the most popular ways to use data in optimization. It is well known that sample average optimization suffers from downward bias. We propose to use smooth estimators rather than empirical ones in optimization problems. We establish consistency results for the optimal value and the set of optimal solutions of the new problem formulation. We analyze the bias of the new problems and identify sufficient conditions for ensuring less biased estimation of the optimal value of the true problem. We show that those conditions are satisfied for many popular statistical problems such as regression models, classification problems, and optimization problems with Average (Conditional) Value-at-Risk. Our numerical experience shows that the new estimators frequently exhibit also smaller variance and smaller mean-square error than those of SAA.

This is a joint work with Yang Lin, Stevens Institute of Technology.

6:30PM – 7:00PM – Jun-ya Gotoh, Chuo University, Japan (jgoto@indsys.chuo-u.ac.jp)

Title: Worst-case Sensitivity

Abstract: We introduce the notion of Worst-Case Sensitivity, defined as the worst-case rate of increase in the expected cost of a Distributionally Robust Optimization (DRO) model when the size of the uncertainty set vanishes. We show that worst-case sensitivity is a generalized measure of deviation and that a large class of DRO models are essentially mean-(worst-case) sensitivity problems when uncertainty sets are small, unifying recent results on the relationship between DRO and regularized empirical optimization with worst-case sensitivity playing the role of the regularizer. More generally, DRO solutions can be sensitive to the family and size of the uncertainty set, and reflect the properties of its worst-case sensitivity. We derive closed-form expressions of worst-case sensitivity for well known uncertainty sets including smooth phi-divergence, total variation, ``budgeted" uncertainty sets, uncertainty sets corresponding to a convex combination of expected value and CVaR, and the Wasserstein metric. These can be used to select the uncertainty set and its size for a given application. This is a joint work with Michael J. Kim (UBC, Canada) and Andrew E.B. Lim (NUS, Singapore).

Uncertainty Management and Machine Learning in Engineering Applications

November 17, 2020
Speakers

10:00AM – 10:30AM – Harbir Antil, George Mason University (hantil@gmu.edu)

Title: Role of Fractional DNNs in Inverse Problems

Abstract: The goal of this talk is to introduce novel fractional derivative based Deep Neural Networks (DNNs). Fractional derivatives have less smoothness requirements and they also have the distinct ability to account for iteration history. The latter enables faster and more accurate training by overcoming the vanishing gradient problem.

We plan to discuss the approximation properties of DNNs and will apply the proposed DNNs to image denoising and tomographic reconstruction problems. We shall establish that they are also excellent surrogates to PDEs and inverse problems with multiple advantages over the traditional methods. We plan to conclude the talk by showing some of our initial results on chemically reacting flows using DNNs.

10:30AM – 11:00AM – Bogdan Grechuk, University of Leicester, England (bg83@leicester.ac.uk)

Title: Mathematical Foundations for Error Correction in Machine Learning

Abstract: All machine learning (ML) and artificial Intelligence (AI) systems make errors. These errors should be corrected without damage of existing skills and avoiding direct human expertise. Recently, a new error-correction mechanism for ML and AI systems has been developed. At the heart of this mechanism are new stochastic separation theorems, which state that linear classifiers in their classical Fisher's form are powerful enough to separate errors from correct responses with high probability, even for exponentially large samples. In this work, we derive, for many families of distributions, separation theorems with very strong, often optimal estimates.

11:00AM – 11:30AM – Grigoriy Zrazhevsky, Kiev University (zgrig@univ.kiev.ua)
Stan Uryasev, Stony Brook University (stanislav.Uryasev@stonybrook.edu)

Title: A Wave Dynamics Inverse Problem: Mathematical Analytics vs Machine Learning

Abstract: The paper considers the problem of evaluation of defectiveness of an elastic rod. The model analyzes the shift of natural resonance frequencies. We solved a direct boundary value problem with a new approach based on generalized functions and asymptotic analysis. We solved inverse problems for finding parameters of one defect and overall defectiveness of a rod with many defects. Problems are reduced to inverse problems for finding zeros of a multivalued function of many variables. Standardly used numerical-analytical methods are not applicable because of small parameters that are separated in the fourth approximation term only. We benchmarked the numerical-analytical method with the Bootstrap-aggregated regression trees (combination of RFA and CART). Sampling was used for data generation. Pros and cons of formal analytical and machine learning approaches are analyzed.

11:30AM – 12:00PM – Kostas Spiliopoulos, Boston University (kspiliop@bu.edu)

Title: DGM: A Deep Learning Method to Solve PDEs.

Abstract: High-dimensional PDEs have been a longstanding computational challenge. We propose to solve high-dimensional PDEs by approximating the solution with a deep neural network which is trained to satisfy the differential operator, initial condition, and boundary conditions. Our algorithm is meshfree, which is key since meshes become infeasible in higher dimensions. Instead of forming a mesh, the neural network is trained on batches of randomly sampled time and space points. The algorithm is tested on a class of high-dimensional free boundary PDEs (American Options), which we are able to accurately solve in up to 200 dimensions. The algorithm is also tested on a high-dimensional Hamilton-Jacobi-Bellman PDE and Burgers' equation. The deep learning algorithm approximates the general solution to the Burgers' equation for a continuum of different boundary conditions and physical conditions (which can be viewed as a high-dimensional space). We call the algorithm a "Deep Galerkin Method (DGM)" since it is similar in spirit to Galerkin methods, with the solution approximated by a neural network instead of a linear combination of basis functions. In addition, we prove a theorem regarding the approximation power of neural networks for a class of quasilinear parabolic PDEs.

12:00PM – 12:30PM – Break

12:30PM- 1:00PM - Johannes Royset, Naval Postgraduate School (joroyset@nps.edu)

Title: Diametrical Risk Minimization: Theory and Computations

Abstract: The theoretical and empirical performance of Empirical Risk Minimization (ERM) often suffers when loss functions are poorly behaved with large Lipschitz moduli and spurious sharp minimizers. We propose and analyze a counterpart to ERM called Diametrical Risk Minimization (DRM), which accounts for worst-case empirical risks within neighborhoods in parameter space. In contrast to common robustification strategies based on perturbing the data set and probability distribution, our approach “diametrically" modifies any solution and thereby obtains theoretical stability guarantees even for poorly behaved functions. DRM has generalization bounds that are independent of Lipschitz moduli for convex as well as nonconvex problems and it can be implemented using a practical algorithm based on stochastic gradient descent. Numerical results illustrate the ability of DRM to find quality solutions with low generalization error in chaotic landscapes from benchmark neural network classification problems with corrupted labels.

1:00PM – 1:30PM – Eugene Feinberg, Stony Brook University (eugene.feinberg@stonybrook.edu)
Rui Ding, Stony Brook University (rui.ding.1@stonybrook.edu)

Title: CVaR Optimization for Sequential Decisions Processes

Abstract: We study CVar optimization for Markov decision processes. It is known that this problem can be reduced to a robust optimization problem with an extended state space. The extended states are vector consisting of two components. The first component is the original state. The second component can be interpreted as the assigned risk, which may be unknown to the decision maker. We show how this problem can be solved by the decision maker without knowing exact values of assigned risks.

1:30PM – 2:00PM – Wilkins Aquino, Duke University (wilkins.aquino@duke.edu)

Title: An Adaptive Sample-Based Approximation Approach for Stochastic Inverse Problems

Abstract: In this work, we adopt a framework based on a Gibbs posterior for updating belief distributions for inverse problems governed by PDEs. In contrast with regular Bayesian methods where distributions of noise are assumed to be known exactly, the Gibbs posterior update does not require a likelihood function. Hence, no exact model for the noise is needed. Instead, the Gibbs posterior is applicable where the unknown parameters are connected to the data through a loss function. We employ a sample based discretization to approximate the continuous prior distribution, which after applying the Gibbs update results in an explicit formula for the posterior weights associated with each sample. To control the number of samples and provide efficient approximations to the posterior, we borrow ideas from sequential Monte Carlo methods to adaptively add samples that cluster within the support of the posterior in a sequential manner. To manage the cost of propagating an increasing number of samples through the loss function, we employ a local reduced basis method to build efficient surrogate models. We demonstrate the performance of our approach through several numerical examples.

2:00PM – 2:30PM – Nat Trask, Sandia National Laboratories (natrask@sandia.gov)

Title: A Data-Driven Exterior Calculus for Learning Models with Exact Physics

Abstract: We consider the extraction of models from observations of a physical system, where derivation of a model from first principles is intractable. Recent works seek to incorporate physical biases into the deep learning process to penalize deviations of the neural network from solutions to a given differential operator. This process weakly enforces physical principles by penalty, and thus physics hold only to within optimization error. For many applications, particularly related to electromagnetics and fluid flow, such properties must hold exactly. We present a data-driven exterior calculus using graphs as surrogates for the mathematical structures underpinning mimetic discretizations of PDEs. This framework allows engineering of deep architectures to guarantee exact enforcement of physics without the need to introduce either regularization or equality constraints into training, providing guarantees of solvability and structure preservation in extracted data-driven models.

2:30PM – 3:00PM – Break

3:00PM – 3:30PM – Ahmad Rushdi, Sandia National Laboratories (arushdi@sandia.gov)

Title: Estimating Neural Networks' Predictive Uncertainty in SciML

Abstract: Neural network models have attracted a lot of research attention in Scientific Machine Learning (SciML) problems. However, they tend to be overconfident when reporting typical point-estimate predictions in classification and regression problems. This could be very harmful when dealing with costly numerical simulations or high-stakes decisions in national security applications. In this talk, we examine uncertainty quantification techniques for neural network models. To understand their variability, we rely on different sources of randomness associated with training samples, weight initialization, dropout methods, and ensemble formations. Motivated by typical SciML situations, we assume a limited sample budget, noisy training data, and suggest approaches for reporting and possibly reducing uncertainty.

3:30PM – 4:00PM – Stan Uryasev, Stony Brook University (stanislav.uryasev@stonybrook.edu)

Title: Renyi Entropy and Calibration of Distribution Tails

Abstract: Joint presentation with Bogdan Grechuk, Michael Zabarankin, Alexey Zrazhevsky

The paper discusses new algorithms for fitting distributions. The popular Shannon entropy maximization approach yields exponential distributions with light tails. Therefore, this approach is not suitable for estimation of probability distributions with heavy tails, which arise in various applications, including financial engineering, reliability theory, and climatology. This paper provides a general solution of the Renyi entropy maximization problem subject to moment constraints. Renyi entropy maximization leads to distributions with heavy tails. In particular, Generalized Pareto Distribution (GPD) can be obtained by maximizing the Renyi entropy with a constraint on mean value. The paper proposed a new approach for estimating the shape parameter of GPD based on mean of logarithm of random value. We used quantile and CVaR regression for the estimation of a conditional GPD distribution tail (as a function of explanatory factors). We have conducted a case study validating the proposed methodology and estimated the conditional tail distributions. Quantile and CVaR regressions are done with the Portfolio Safeguard (PSG) optimization package which has precoded Koenker-Bassett and Rocafellar errors for quantile and CVaR regressions. Case study is posted at the web (codes, data, results).

4:00PM – 4:30PM – Pavlo Krokhmal, University of Arizona (krokhmal@email.arizona.edu)

Title: Risk-Averse Graph Theoretical Problems

Abstract: In this talk we discuss risk-averse stochastic optimization in discrete systems, such networks or graphs. In particular, we consider identification of minimum-risk structures in graphs with random vertex or edge weights The risk aversity is achieved by application of a family of certainty equivalent coherent/convex measures of risk, resulting in mixed-integer conic programming with nonsymmetric cones. A connection of this framework to determining the systemic risk of a networked system is discussed. In addition, we discuss the interplay of risk reduction and diversification in discrete systems. An application of the developed concepts to logistics networks is presented.

4:30PM – 5:00PM – Michael Zabarankin, Stevens Institute of Technology (michael.zabarankin@stevens.edu)

Title: Regression Analysis: Likelihood, Error and Entropy

Abstract: In a regression with independent and identically distributed normal residuals, the log-likelihood function yields an empirical form of the L2 norm, whereas the normal distribution can be obtained as a solution of differential entropy maximization subject to a constraint on the L2 norm of a random variable. The L1 norm and the double exponential (Laplace) distribution are related in a similar way. These are examples of an "inter-regenerative" relationship. In fact, L2 norm and L1 norm are just particular cases of general error measures introduced by Rockafellar et al. on a space of random variables. General error measures are not necessarily symmetric with respect to ups and downs of a random variable, which is a desired property in finance applications where gains and losses should be treated differently. This work identifies a set of all error measures, denoted by E, and a set of all probability density functions (PDFs) that form "inter-regenerative" relationships (through log-likelihood and entropy maximization). It also shows that M-estimators, which arise in robust regression but, in general, are not error measures, form "inter-regenerative" relationships with all PDFs. In fact, the set of M-estimators, which are error measures, coincides with E. On the other hand, M-estimators are a particular case of L-estimators that also arise in robust regression. A set of L-estimators which are error measures is identified—it contains E and the so-called trimmed Lp norms.

5:00PM – 5:30PM – Break

5:30PM – 6:00PM – Alex Lipton, Hebrew University of Jerusalem, MIT (alexlipt@mit.edu)

Title: Observations of a financial supernova: Emerging Trends in Decentralized Finance (DeFi)

Abstract: In this talk, we discuss some of the most recent developments in the cryptocurrency ecosystem. Specifically, we review stable coins, their classification, potential applications, and related topics, and closely related emerging field of DeFi (Decentralized Finance), including Automated Market Makers (AMM), yield farming, and other peculiar concepts. We explain mathematics, economics, and technology behind these developments and elaborate on their pros and cons.

6:00PM – 6:30PM – James Ostrowski, University of Tennessee (jostrows@utk.edu)

Title: Quantum Approximate Optimization Algorithm

Abstract: The quantum approximate optimization algorithm (QAOA) is a method of approximately solving combinatorial optimization problems. While QAOA is developed to solve a broad class of combinatorial optimization problems, it is not clear which classes of problems are best suited for it. We will discuss a variety of strategies, grounded in polyhedral theory, for attempting to identify which combinatorial optimization problems are likely to display quantum dominance (if any).

6:30PM – 7:00PM – Gianluca Iaccarino, Stanford University (jops@stanford.edu)

Title: Data-free and Data-driven Uncertainty Quantification in Turbulence Simulations

Abstract: Despite continued advances in high-fidelity turbulent flow simulations, closure models based on the Reynolds-Averaged Navier-Stokes (RANS) equations are projected to remain in use for considerable time to come. However, it is common knowledge that RANS predictions are corrupted by epistemic model-form uncertainty to a degree which is unknown a-priori. Hence, to obtain a computational framework of predictive utility, a model-form Uncertainty Quantification framework is indispensable. Applying the spectral decomposition to the modeled Reynolds-Stress Tensor (RST) allows for the introduction of decoupled perturbations into the baseline intensity (kinetic energy), shape (eigenvalues), and orientation (eigenvectors). Within this perturbation framework, we look for a-priori known limiting physical bounds. Since these bounds are universal, they can be used to constrain uncertainty estimates in any predictive flow scenario. Thus, even in the absence of relevant training data, we can maximize the spectral perturbations in order to obtain conservative uncertainty intervals. On the other hand, any high-fidelity reference data can be used to further constrain the uncertainty estimates using commonly available data assimilation techniques. We will demonstrate our framework on canonical flow problems using random forest regression to incorporate. DNS data into the uncertainty estimates of conventional RANs closures. modeling, and pricing theory. Time permitting we will cover connections to machine learning.