samsi
NSF·Duke·NCSU·UNC

QMC Opening Workshop
August 28-September 1, 2017

SPEAKER TITLES/ABSTRACTS

Jose Blanchet
Stanford University

"New Problems and Algorithms at the Interface of Optimal Transport, Statistics and Operations Research"

Recently, there has been a surge in activity at the interface of optimal transport and statistics (with special emphasis on machine learning applications). The talk will summarize new results and challenges in this active area. For example, we will show how many of the most popular estimators in machine learning (such as Lasso and svm's) can be interpreted as games. This interpretation opens the door for new and potentially better estimators and algorithms, as well as questions about the underlying complexity of these new class of estimators.

(This talk is based on joint work with F. He, Y. Kang, K. Murthy, and F. Zhang)

Mahadevan Ganesh
Colorado School of Mines

"A Sign-definite Heterogeneous Media Wave Propagation Model"

The standard Galerkin formulation of the acoustic wave propagation, governed by the Helmholtz partial differential equation (PDE), is indefinite for large wavenumbers. However, the Helmholtz PDE is in general not indefinite. The lack of coercivity (indefiniteness) is one of the major difficulties for approximation and simulation of heterogeneous media wave propagation models, including application to stochastic wave propagation Quasi Monte Carlo (QMC) analysis. We will present a new class of sign-definite continuous and discrete preconditioned FEM Helmholtz wave propagation models.

Mathieu Gerber
University of Bristol

"Introduction to Sequential quasi-Monte Carlo"

Sequential quasi-Monte Carlo (SQMC) is a quasi-Monte Carlo (QMC) version of sequential Monte Carlo (or particle filtering), a popular class of Monte Carlo techniques used to carry out inference in state space models. In this talk I will first review the SQMC methodology as well as some theoretical results. Although SQMC converges faster than the usual Monte Carlo error rate its performance deteriorates quickly as the dimension of the hidden variable increases. However, I will show with an example that SQMC may perform well for some "high" dimensional problems. I will conclude this talk with some open problems and potential applications of SQMC in complicated settings.

Mike Giles
Oxford University

"QMC and Thinning for Empirical Datasets"

In this talk we consider the question of how to use QMC with an empirical dataset, such as a set of points generated by MCMC. Using ideas from partitioning for parallel computing, we apply recursive bisection to reorder the points, and then interleave the bits of the QMC coordinates to select the appropriate point from the dataset. Numerical tests show that in the case of known distributions this is almost as effective as applying QMC directly to the original distribution.

The same recursive bisection can also be used to thin the dataset, by recursively bisecting down to many small subsets of points, and then randomly selecting one point from each subset. This makes it possible to reduce the size of the dataset greatly without significantly increasing the overall error.

Co-author: Fei Xie

Fred Hickernell
Illinois Institute of Technology

"Error Analysis for Quasi-Monte Carlo Methods"

Multidimensional integrals may be approximated by weighted averages of integrand values. Quasi-Monte Carlo (QMC) methods are more accurate than simple Monte Carlo methods because they carefully choose where to evaluate the integrand. This tutorial focuses on how quickly QMC methods converge to the correct answer as the number of integrand values increases. The answer may depend on the smoothness of the integrand and the sophistication of the QMC method. QMC error analysis may assumes the integrand belongs to a reproducing kernel Hilbert space or may assume that the integrand is an instance of a stochastic process with known covariance structure. These two approaches have interesting parallels. This tutorial also explores how the computational cost of achieving a good approximation to the integral depends on the dimension of the domain of the integrand. Finally, this tutorial explores methods for determining how many integrand values are needed to satisfy the error tolerance. Relevant software is described.

James M. Hyman
Tulane University,

"High Accuracy Algorithms for Interpolating and Integrating Multivariate Functions Defined by Sparse Samples in High Dimensions"

We will describe and analyze accurate and efficient numerical algorithms to interpolate and approximate the integral of multivariate functions. The algorithms can be applied when we are given the function values at an arbitrary positioned, and usually small, existing sparse set of function values (samples), and additional samples are impossible, or difficult (e.g. expensive) to obtain. The methods are based on local, and global, tensor-product sparse quasi-interpolation methods that are exact for a class of sparse multivariate orthogonal polynomials.
Although the algorithms can be applied for general distributions of sample points, they are especially effective in improving the approximation of integrals defined on low-discrepancy (approximately uniformly distributed) sample. One advantage of this approach is that it continues to be effective on even high-discrepancy sampling distributions. This allows the algorithms to be used with adaptive sampling algorithms that concentrate sample points in regions with locally high gradients.

We will compare the convergence rate of the method on smooth and discontinuous functions defined on both low discrepancy quasi-Monte Carlo distributions of samples, as well as for sample distributions that are far from uniformly distributed. We observe that although the approach does not change the error convergence rate, the constant multiplicative factor in the error is often reduced by a factor of over 100 or more.

Joint research with Jeremy Dewar, Lin Li, and Mu Tian

Peter Kritzer
Austrian Academy of Sciences

"Numerical Integration in Hermite Spaces"

In this talk, we give an overview of results on numerical integration in Hermite spaces. These spaces contain functions defined on $\mathbb{R}^d$, and can be characterized by the decay of their Hermite coefficients. We consider the case of exponentially as well as polynomially decaying Hermite coefficients. For numerical integration, we either use Gauss-Hermite quadrature rules or algorithms based on quasi-Monte Carlo rules. We present upper and lower error bounds for these algorithms, and discuss their dependence on the dimension d. Furthermore, we comment on open problems for future research.

Frances Kuo
UNSW Sydney

"Application of QMC to PDEs with Random Coefficients -- a survey of analysis and implementation"

In this tutorial I will provide a survey of recent research efforts on the application of QMC methods to PDEs with random coefficients. Such PDE problems occur in the area of uncertainty quantification. A prime example is the flow of water through a disordered porous medium. There is a huge body of literature on this topic using a variety of methods. QMC methods are relatively new to this application area. The aim of this tutorial is to provide an entry point for QMC experts wanting to start research in this direction, for PDE analysts and practitioners wanting to tap into contemporary QMC theory and methods, and for any one else who sees how to cross-fertilize the ideas to other application areas.

Pierre L'Ecuyer
University of Montreal

"Lattice Rules for Quasi-Monte Carlo"

Lattice rules are one of the two main classes of methods for quasi-Monte Carlo (QMC) and randomized quasi-Monte Carlo (RQMC) integration. In this tutorial, we recall the definition and summarize the key properties of lattice rules. We discuss what classes of functions these rules are good to integrate, and how their parameters can be chosen in terms of variance bounds for these classes of functions. We consider integration lattices in the real space as well as in a polynomial space over the finite field F2. We provide various numerical examples of how these rules perform compared with standard Monte Carlo. Some examples involve high-dimensional integrals, others involve Markov chains. We also discuss software design for RQMC and what software is available.

Lester Mackey
Stanford University

"Measuring Sample Quality with Stein's Method"

To improve the efficiency of Monte Carlo estimation, practitioners are turning to biased Markov chain Monte Carlo procedures that trade off asymptotic exactness for computational speed. The reasoning is sound: a reduction in variance due to more rapid sampling can outweigh the bias introduced. However, the inexactness creates new challenges for sampler and parameter selection, since standard measures of sample quality like effective sample size do not account for asymptotic bias. To address these challenges, we introduce new computable quality measures based on Stein's method that quantify the maximum discrepancy between sample and target expectations over a large class of test functions. We use our tools to compare exact, biased, and deterministic sample sequences and illustrate applications to hyperparameter selection, biased sampler selection, one-sample hypothesis testing, and sample quality improvement.

Simon Mak
Georgia Tech University

“Support Points - a new way to compact distributions, with small-data and big-data applications”

This talk introduces a new way to compact a (possibly non-uniform) probability distribution F into a set of representative points, called *support points*. These point sets can have important uses for both small-data problems, such as experimental design and uncertainty quantification in engineering applications, as well as big-data problems, such as the optimal reduction of large datasets in Bayesian computation. We first present support points as the minimizer of a powerful goodness-of-fit test called the energy distance, and discuss why such point sets are appealing to use for simulation and integration. An extension of this point set, called *projected support points*, is then introduced for high-dimensional integration under non-uniform F. We show that support points (and its variants) can provide good solutions to the aforementioned small-data and big-data problems. This talk concludes with some new ideas and ongoing work on experimental design, potential theory and robust optimization.

Dirk Nuyens
KU Leuven

“Higher-order Convergence for Integration on R^N”

We study the application of higher-order cubature rules to approximate infinite-variate integrals over R^N. We analyse the dimension-truncation error and the cubature rule error for each finite set u in an MDM decomposition and look for conditions in the form of POD weights to achieve the total error to decay at an algebraic higher order rate in terms of the invested work.
This is joint work with Dong T.P. Nguyen (KU Leuven, Belgium)
Chris Oates
Newcastle University

“Bayesian Probabilistic Numerical Methods (Part I)”

In this work, numerical computation - such as numerical solution of a PDE - is treated as an inverse problem in its own right. The popular Bayesian approach to inversion is considered, wherein a posterior distribution is induced over the object of interest by conditioning a prior distribution on the same finite information that would be used in a classical numerical method. The main technical consideration is that the data in this context are non-random and thus the standard Bayes' theorem does not hold. General conditions will be presented under which such Bayesian probabilistic numerical methods are well-posed, and a sequential Monte-Carlo method will be shown to provide consistent estimation of the posterior. The paradigm is extended to computational ``pipelines", through which a distributional quantification of numerical error can be propagated. A sufficient condition is presented for when such propagation can be endowed with a globally coherent Bayesian interpretation, based on a novel class of probabilistic graphical models designed to represent a computational work-flow. The concepts are illustrated through explicit numerical experiments involving both linear and non-linear PDE models. Full details are available in arXiv:1702.03673.

Art Owen
Stanford University

“Introduction to Quasi-Monte Carlo”

A fundamental numerical problem in many sciences is to compute integrals. These integrals can often be expressed as expectations and then approximated by sampling methods. Monte Carlo sampling is very competitive in high dimensions, but has a slow rate of convergence. One reason for this slowness is that the MC points form clusters and gaps. Quasi-Monte Carlo methods greatly reduce such clusters and gaps, and under modest smoothness demands on the integrand they can greatly improve accuracy. This can even take place in problems of surprisingly high dimension. This talk will introduce the basics of QMC and randomized QMC. It will include discrepancy and the Koksma-Hlawka inequality, some digital constructions and some randomized QMC methods that allow error estimation and sometimes bring improved accuracy.

Clémentine Prieur
Université Grenoble Alpes

“Introduction to Global Sensitivity”

Many mathematical models use a large number of poorly-known parameters as inputs. Quantifying the influence of each of these parameters is one of the aims of sensitivity analysis. Global Sensitivity Analysis is an important paradigm for understanding model behavior, characterizing uncertainty, improving model calibration, etc. Inputs’ uncertainty is modeled by a probability distribution. There exist various measures built in that paradigm. This tutorial focuses on the so-called Sobol’ indices, based on functional variance analysis. Estimation procedures will be presented, and the choice of the designs of experiments these procedures are based on will be discussed. As Sobol’ indices have no clear interpretation in the presence of statistical dependences between inputs, it also seems promising to measure sensitivity with Shapley effects, based on the notion of Shapley value, which is a solution concept in cooperative game theory.

Florian Puchhammer
University of Montreal

“Lower Bounds for the Discrepancy of Point Sets and Sequences”

QMC algorithms usually rely on a choice of N evenly distributed integration nodes in $[0,1)^d$. A common means to assess such an equidistributional property for a point set or sequence is the so-called discrepancy function, which compares the actual number of points to the expected number of points (assuming uniform distribution on $[0,1)^{d}$) that lie within an arbitrary axis parallel rectangle anchored at the origin. The dependence of the integration error using QMC rules on various norms of the discrepancy function is made precise within the well-known Koksma--Hlawka inequality and its variations. In many cases, such as L^{p} spaces, $1<p<\infty$, the best growth rate in terms of the number of points N as well as corresponding explicit constructions are known. In the classical setting $p=\infty$ sharp results are absent for $d\geq3$ already and appear to be intriguingly hard to obtain. This talk shall serve as a survey on discrepancy theory with a special emphasis on the L^{∞} setting. Furthermore, it highlights the evolution of recent techniques and presents the latest results.

Christoph Schwab
SAM, ETH Zurich

"Multilevel QMC for Forward and Inverse UQ"

We present recent result on the numerical analysis of Quasi Monte-Carlo quadrature methods, applied to forward and inverse uncertainty quantification for elliptic and parabolic PDEs. Particular attention will be placed on Higher-Order QMC, the stable and efficient generation of interlaced polynomial lattice rules, and the numerical analysis of multilevel QMC Finite Element discretizations with applications to computational uncertainty quantification.

Ian H. Sloan
University of New South Wales

"Generating Random Fields the Circulant Way"

The generation of Gaussian random fields over a physical domain is a challenging problem in computational mathematics, especially when the correlation length is short and the field is rough. The traditional approach is to make use of a truncated Karhunen-Loeve (KL) expansion, but the generation of even a single realisation of the field may then be effectively beyond reach (especially for 3-dimensional domains) if the need is to obtain an expected L2 error of say 5%, because of the potentially very slow convergence of the KL expansion. In this talk, based on joint work with Ivan Graham, Frances Kuo, Dirk Nuyens, and Rob Scheichl, a completely different approach is used, in which the field is initially generated at a regular grid on a 2- or 3-dimensional rectangle that contains the physical domain, and then possibly interpolated to obtain the field at other points. In that case there is no need for any truncation. Rather the main problem becomes the factorisation of a large dense matrix. For this we use circulant embedding and FFT ideas. Quasi-Monte Carlo integration is then used to evaluate the expected value of some functional of the finite-element solution of an elliptic PDE with a random field as input.

Tim Sullivan
Free University of Berlin/Zuse Institute Berlin

“Bayesian Probabilistic Numerical Methods (Part II)”

(This is the second part of joint talk with Chris Oates. Time permitting, the two talks will address the theoretical and computational aspects respectively.)

In this work, numerical computation - such as numerical solution of a PDE - is treated as an inverse problem in its own right. The popular Bayesian approach to inversion is considered, wherein a posterior distribution is induced over the object of interest by conditioning a prior distribution on the same finite information that would be used in a classical numerical method. The main technical consideration is that the data in this context are non-random and thus the standard Bayes' theorem does not hold. General conditions will be presented under which such Bayesian probabilistic numerical methods are well-posed, and a sequential Monte-Carlo method will be shown to provide consistent estimation of the posterior. The paradigm is extended to computational ``pipelines", through which a distributional quantification of numerical error can be propagated. A sufficient condition is presented for when such propagation can be endowed with a globally coherent Bayesian interpretation, based on a novel class of probabilistic graphical models designed to represent a computational work-flow. The concepts are illustrated through explicit numerical experiments involving both linear and non-linear PDE models. Full details are available in arXiv:1702.03673.

Roshan Vengazhiyil
Georgia Institute of Technology

“Deterministic Sampling for Bayesian Computation”

Markov chain Monte Carlo (MCMC) methods are popularly used in Bayesian computation. However, they need large number of samples for convergence which can become costly when the posterior distribution is expensive to evaluate. Deterministic sampling techniques such as Quasi-Monte Carlo (QMC) can be a useful alternative to MCMC, but the existing QMC methods are mainly developed only for sampling from unit hypercubes. Unfortunately, the posterior distributions can be highly correlated and nonlinear making them occupy very little space inside a hypercube. Thus, most of the samples from QMC can get wasted. The QMC samples can be saved if they can be pulled towards the high probability regions of the posterior distribution using inverse probability transforms. But this can be done only when the distribution function is known, which is rarely the case in Bayesian problems. In this talk, I will discuss a deterministic sampling technique, known as minimum energy designs, which can directly sample from the posterior distributions.

G. W. Wasilkowski
University of Kentucky

"∞-Variate Integration"

We present results on efficient approximation of integrals with infinitely many variables. We provide new concepts of worst case truncation and superposition dimensions and show that, under modest error demands, these dimensions are very small for functions from weighted tensor product spaces. We also present Multivariate Decomposition Method that is almost as efficient as quadratures for univariate problems.

The presentation is based on papers co-authored with A. Gilbert, M. Gnewuch, M. Hefter, P. Kritzer, F. Y. Kuo, F. Pillichshammer, L. Plaskota, K. Ritter, I. H. Sloan, and H. Wozniakowski.

Clayton Webster
Oak Ridge National Laboratory

"Sparse Polynomial Approximation via Compressed Sensing of High Dimension Functions"

In this talk, we present a compressed sensing approach to polynomial approximation of functions in high dimensions. Of particular interest is the parameterized PDE setting, where the target function is smooth, characterized by a rapidly decaying orthonormal expansion, whose most important terms are captured by a lower (or downward closed) set. By exploiting this fact, we develop a novel weighted minimization procedure with a precise choice of weights, and a modification of the iterative hard thresholding method, for imposing the downward closed preference. Moreover, the recovery of the corresponding best approximation using our methods is established through an improved bound for the restricted isometry property. In addition, we will also present a new theory revealing that nonconvex minimizations are at least as good as ℓ_1 minimization in exact recovery of sparse signals. These theoretical recovery guarantees are developed through a unified null space property based-condition that encompasses all currently proposed nonconvex functionals in literature. Numerical examples are provided to support the theoretical results and demonstrate the computational efficiency of our new weighted approach, as well as nonconvex minimizations.

Henryk Wozniakowski
Columbia University

“Quasi-polynomial Tractability of Linear Tensor Products using Function Values”

We study QPT (quasi-polynomial tractability) in the worst case setting of linear tensor product problems defined over Hilbert spaces. We prove QPT for algorithms that use only function values under three assumptions'

1. the minimal errors for the univariate case decay polynomially fast to zero,
2. the largest singular value for the univariate case is simple,
3. the eigenfunction corresponding to the largest singular value is a multiple of the function value at some point.

The first two assumptions are necessary for QPT. The third assumption is necessary for QPT for some Hilbert spaces. Joint work with Erich Novak

Dongbin Xiu
Ohio State University

“Sequential Function Approximation in High Dimensions with Big Data”

One of the central tasks in computational mathematics and statistics is to accurately approximate unknown target functions. This is typically done with the help of data â€” samples of the unknown functions. The emergence of Big Data presents both opportunities and challenges. On one hand, big data introduces more information about the unknowns and, in principle, allows us to create more accurate models. On the other hand, data storage and processing become highly challenging. In this talk, we present a set of sequential algorithms for function approximation in high dimensions with large data sets. The algorithms are of iterative nature and involve only vector operations. They use one data sample at each step and can handle dynamic/stream data. We present both the numerical algorithms, which are easy to implement, as well as rigorous analysis for their theoretical foundation.

Guannan Zhang
Oak Ridge National Laboratory

“Probabilistic Numerical Methods for High-Dimensional Partial Integral Differential Equations”

In this talk, we discuss some recent advances in probabilistic schemes for high-dimensional PIDEs. It is known that traditional PDE solvers, e.g., finite element, finite difference methods, do not scale well with the increase of dimension. The idea of probabilistic schemes is to link a wide class of nonlinear parabolic PIDEs to stochastic Levy processes based on nonlinear version of the Feynman-Kac theory. As such, the solution of the PIDE can be represented by a conditional expectation (i.e., a high-dimensional integral) with respect to a stochastic dynamical system driven by Levy processes. In other words, we can solve the PIDEs by performing high-dimensional numerical integration. A variety of quadrature methods could be applied, including MC, QMC, sparse grids, etc. The probabilistic schemes have been used in many application problems, e.g., particle transport in plasmas (e.g., Vlasov-Fokker-Planck equations), nonlinear filtering (e.g., Zakai equations), and option pricing, etc.