Submittee: Chen Greif
Date Submitted: 2013-09-11 11:57
Title: Workshop on Numerical Linear Algebra and Optimization
Event Type: Conference-Workshop

Location:
UBC, ESB 2012

Dates:
August 8-10, 2013

Topic:
Optimization problems involving eigenvalues, including issues related to semidefinite programming, nonconvex problems involving constraints on the eigenvalues of nonsymmetric matrices, or their generalization, namely pseudospectra, applications of eigenvalue optimization to fields such as control theory and machine learning, the efficient use of tools from numerical linear algebra for large-scale eigenvalue optimization.

Methodology:
The workshop was comprised of plenary lectures (one stream), and a poster session in the first evening. Each session concerned (primarily) either numerical linear algebra or optimization, featuring two or three plenary lectures. In each day there was an approximately even split between the numerical linear algebra and optimization talks, highlighting numerical linear algebra and optimization aspects of relevant topics. There was a good match among topics within each session. For example, a session was devoted to non-smooth analysis with eigenvalue optimization and applications in mind, featuring Adrian Lewis, James Burke and Frank Curtis. Another session featuring Volker Mehrmann, Didier Henrion and Sara Grundel focused on problems in control theory relevant to eigenvalue optimization. The theme of another session featuring Michael Saunders, Daniel Szyld and Josef Sifuentes was focused on Krylov subspace methods, which constitute indispensable numerical ingredients for optimization algorithms, including eigenvalue optimization.

Objectives Achieved:
Eigenvalue optimization problems in the convex and non-convex settings possess intrinsic similarities. Yet they are typically studied by different communities: convex problems, especially those that can be formulated as SDPs, have received immense attention in optimization, while non-convex problems, such as those concerning pseudospectra, have been hot research directions in numerical linear algebra. // The workshop brought these two communities together and provided the opportunity to exchange ideas. On the first day, the opening was made with a session partly related to matrices with multiple eigenvalues, in particular distances to such matrices, that are of interest to the numerical linear algebra community. Various other talks, for instance the talks by Shreemayee Bora and Daniel Kressner, focused on similar distance problems popular in numerical linear algebra, leading to non-convex eigenvalue optimization problems. On the side of the optimization community, quite a few talks concerned SDP relaxations of NP-hard problems, e.g., the talks by Henry Wolkowicz, Stephen Vavasis and Franz Rendl. Thus the workshop provided an

opportunity for each of the two communities to become better acquainted with the problems that are of interest to the other community. // Modern applications of optimization problems involving eigenvalues have been shared throughout the workshop. The talk by James Burke elaborated on applications to machine learning, the talks by Volker Mehrmann and Didier Henrion illustrated applications to control theory, the talk by Franz Rendl considered a classical application from graph theory. // Effective use of large-scale numerical linear algebra tools was also a theme of the workshop. An important tool from numerical linear algebra, Krylov subspace methods, was covered by a session consisting of lectures by Michael Saunders, Daniel Szyld and Josef Sifuentes. The use of appropriate matrix factorizations was highlighted by the lectures by Paul Van Dooren and Charles Van Loan. The lecture by Jim Demmel focused on efficient parallel implementation of numerical linear algebra routines.

Scientific Highlights:

The workshop was a huge success in terms of attendance (the room was full and in fact registration had to be closed at some point to avoid room overflow), and the main goal of bringing together scientists from various communities and disciplines was accomplished. There was a lively interaction, and a wealth of fruitful discussions. The workshop featured 27 speakers, who are leading experts in their fields. The talks were of a very high quality, and many of them introduced issues that are at the forefront of numerical linear algebra and optimization research. // A few highlights in terms of talks: 1) Jim Demmel from Berkeley gave an overview of communication-avoiding algorithms for linear algebra, both iterative and direct methods. He presented algorithms with provable lower bound on communications. These algorithms clearly push the envelope in terms of speedup compared to the traditionally used algorithm. // 2) Anne Greenbaum from the University of Washington provided an overview of leading-edge analysis that aims at closing in on a bound on the norm of analytic functions of matrices, to obtain an upper bound stated in Crouzeixâ€™s Conjecture. This is an open problem. // 3) Didier Henrion gave a beautiful presentation that explained how to use numerical methods and symbolic methods to compute linear combinations of matrices whose spectrum is prescribed. This is an important problem in control. // 4) Michael Friendlander from UBC introduced a new framework for duality, that goes beyond the traditional Lagrange-multiplier formulation, and is widely applicable to a very large range of nonsmooth optimization problems. // 5) Michael Saunders gave an up-to-date and fascinating account of the very latest in tridiagonalization techniques in the context of modern Krylov subspace solvers for solving linear systems and eigenvalue problems. // 6) Paul Van Dooren presented a very new technique based on anti-triangular factorization, for solving indefinite problems. The method is more stable than traditional methods that are currently used, and it shows significant promise. // 7) Volker Mehrmann's lecture was on the distance from a stable descriptor system to a nearest unstable one, leading to non-convex eigenvalue optimization problems. Problems of this sort are endorsed by the robust control theory community. // The above are just a small sample of a long list of talks that illustrated where the fields of numerical linear algebra and optimization currently are. The interaction among the workshop participants demonstrated the importance of this type of interaction, and we have every reason to believe that new collaborations and new directions are going to arise from the workshop in the near future.

Organizers:

Ioana Dumitriu (Washington, USA), Chen Greif (UBC, Canada), Emre Mengi (Koc, Turkey)

Speakers:

[Anne Greenbaum's talk is a attached as a PDF file, due to the many symbols the abstract contains; all other talks are listed below.] /// Shreemayee Bora, IIT Guwhati // Title: Structured eigenvalue backward errors for structured matrix polynomials. // Abstract: We consider the problem of obtaining the backward error of $\lambda \in \C$ considered as an approximate eigenvalue of a structured matrix polynomial with respect to structure preserving perturbations and present formulas for these backward errors for certain structures. Numerical examples suggest that in many instances there is a considerable difference in the values of the backward errors with respect to structure preserving and arbitrary perturbations. (Joint work with Christian Mehl and Michael Karow of TU Berlin and Punit Sharma of IIT Guwahati) /// James Burke, University of Washington // Title: Piecewise linear quadratic and quadratic support functions in regularized regression, machine learning, system identification, and Kalman smoothing. // Abstract: We discuss the class of piecewise linear quadratic (PLQ) penalties. Well known examples include the L2, L1, Huber, Vapnik, hinge loss, elastic net, and many others. These functions play a crucial role in a variety of applications, including machine learning, robust statistical inference, sparsity promotion, and inverse problems such as Kalman smoothing. // We build on a dual representation of this class of penalties, and use it to characterize conditions necessary to interpret these functions as negative logs of true probability densities. We then present a generalized function class called quadratic support (QS) functions that shares this statistical interpretation. // In the second part of the talk, we discuss a general solver for the PLQ class. The dual representation of these penalties allows a simple calculus, which we exploit to build an overall objective from simple parts. The resulting algorithm can handle a number of fairly general problems, while still efficiently exploiting structure. Moreover, simple constraints are easily incorporated. We present several numerical examples from different areas, and end with a discussion of an application of these ideas to Kalman smoothing. (Joint work with Sasha Aravkin, Brad Bell, and Gianluigi Pillonetto.) /// Frank Curtis, Lehigh University // Title: A quasi-Newton gradient sampling algorithm for nonsmooth optimization // Abstract: The gradient sampling algorithm proposed by Burke, Lewis, and Overton provides a theoretically strong framework for solving nonconvex, nonsmooth optimization problems. In this talk, we discuss practical enhancements to the algorithm (adaptive sampling, quasi-Newton Hessian approximations, etc.) that make it more computationally appealing, and discuss extensions for solving constrained problems. Numerical experiments with an algorithm that automatically transitions from a BFGS method to our adaptive gradient sampling scheme show that the method is quite promising as a general-purpose solver for nonconvex, nonsmooth optimization. /// James Demmel, University of California at Berkeley // Title: Communication-avoiding algorithms for linear algebra and beyond // Abstract: Algorithms have two costs: arithmetic and communication, i.e. moving data between levels of a memory hierarchy or processors over a network. Communication costs (measured in time or energy per operation) already greatly exceed arithmetic costs, and the gap is growing over time following technological trends. Thus our goal is to design algorithms that minimize communication. We present algorithms that attain provable lower bounds on communication, and show large speedups compared to their conventional counterparts. These algorithms are for direct and iterative linear algebra, for dense and sparse matrices, as well as direct n-body simulations. Several of these algorithms exhibit perfect strong scaling, in both time and energy: run time (resp. energy) for a fixed problem size drops proportionally to p (resp. is independent of p). Finally, we describe extensions to algorithms involving arbitrary loop nests and array accesses, assuming only that array subscripts are linear functions of the loop indices. /// Maryam Fazel, University of Washington // Title: Recovery of simultaneously structured models from limited observations // Abstract: Recovering a signal with a low-dimensional structure given a limited number of observations is a central problem in signal processing (compressed sensing), machine learning (recommender systems), and system identification. Structured models are often represented by sparse vectors, low-rank matrices, the sum of sparse and low-rank matrices, etc. Existing results characterize the number of observations for successful recovery for these structures. // In many applications, however, the desired model has \emph{multiple} structures simultaneously. Applications include sparse phase retrieval, and learning models with several structural priors in machine learning tasks. Often, penalties that promote individual structures are known (e.g.,ℓ_1 norm for sparsity, nuclear norm for matrix rank), and require a minimal number of generic measurements, so it is reasonable to minimize a combination of such norms. We show that, surprisingly, if we use multiobjective optimization with the individual norms, we can do no better (order-wise) than an algorithm that exploits only one of the structures. This result holds in a general setting and suggests that to fully exploit the multiple structures, we need a new convex relaxation. /// Michael Friedlander, University of British Columbia // Title: A dual approach to sparse optimization // Abstract: A feature common to many sparse optimization problems is that the number of variables may be significantly larger than the number of constraints---e.g., the standard matrix-lifting approach for binary optimization results in a problem where the number of variables is quadratic in the number of constraints. We consider a duality framework applicable to a wide range of nonsmooth sparse optimization problems that allows us to leverage the relatively small number of constraints. Preliminary numerical results illustrate our approach and its flexibility. (Joint work with Nathan Krislock and Ives Macedo.) /// Donald Goldfarb, Columbia University // Title: Low-rank Tensor Recovery: theory and algorithms // /// Abstract: Recovering a low-rank tensor from incomplete or corrupted observations is a recurring problem in signal processing and machine learning. To exploit the structure of data that is intrinsically more than three-dimensional, convex models such low-rank completion and robust principal component analysis (RPCA) have been extended to tensors. In this work, we rigorously establish the recovery guarantees for both tensor completion and tensor RPCA. We also demonstrate that using the most popular convex relaxation for the tensor Tucker rank can be substantially suboptimal in terms of the number of observations needed for exact recovery. We introduce a very simple, new convex relaxation which is shown, both theoretically and empirically, to be greatly superior to the previous model. Moreover, we propose algorithms to solve these low-rank tensor recovery models based on the Accelerated Linearized Bregman (ALB) method and the Alternating Direction Augmented Lagrangian (ADAL) method. Finally, we empirically investigate the recoverability properties of the convex models, and compare the computational performance of the algorithms on both simulated and real data. (Joint work with C. Mu, B. Huang and J. Wright.) /// Anne Greenbaum, University of Washington // Title: K-spectral sets and applications // Abstract: See the attached pdf file (Joint work with Daeshik Choi.) /// Sara Grundel, Max Planck Institute - Magdeburg // Title: Interpolation and classical model order reduction to create parametric model order reduction // Abstract: Model Order Reduction Methods for linear systems are well studied and many successful methods exist. We will review some and explain more recent advances in Parametric Model Order Reduction. The focus will be on methods where we interpolate certain significant measures, that are computed for specific values of the parameters by Radial Basis Function Interpolation. These measures have a disadvantage as they behave like eigenvalues of matrices depending on parameters and we will explain how that can be dealt with in practice. We will furthermore need to introduce a technique to create a medium size model by an extension of barycentric interpolation to make our algorithm efficient. /// Nicola Guglielmi, University of L'Aquila // Title: Differential equations for the approximation of the closest defective matrix // Abstract: Let A be a matrix with distinct eigenvalues. We address the problem of determining the closest defective matrix to A. // We propose a new method for the approximation of the complex and real distance to the set of defective matrices, which is based on a differential equation in the manifold of rank-2 (for the complex case) and rank-4 (in the real case) matrices. An approximation of the closest defective matrix of the form A+E is constructed explicitly with E of low rank. // Due to the local behaviour of the method we are able to compute upper bounds of the distances; nevertheless the algorithms typically find good approximations in cases where we can test this. // Extensions to different structures will also be discussed. (Joint work with P. Butta' and S. Noschese of Universita di Roma La Sapienza and M. Manetta of UniversitÃ dell'Aquila.) /// Didier Henrion, University of Toulouse // Title: Solving structured eigenvalue assignment problems // Abstract: We address the problem of finding a structured control law assigning given closed-loop poles for a linear feedback system. Mathematically, the problem can be formulated as follows: given a set of distinct negative numbers $S=\{s_1,\ldots,s_n\}$ and given n-by-n real matrices $A_0,A_1,\ldots,A_n$, find real numbers $k_1,\ldots,k_n$ such that the matrix $A_0+k_1 A_1+\cdots+k_n A_n$ has spectrum S. We propose two approaches to this problem: (1) a numerical method based on polynomial optimization and semidefinite programming relaxations; (2) a symbolic method based on techniques of computer algebra (Grobner bases, rational univariate representation). We discuss the relative merits and shortcomings of these two approaches applied to a challenging benchmark problem of electrical network design. (Joint work with Sergio Galeani of Roma, Alain Jacquemard of Dijon and Vienna, and Luca Zaccarian ofToulouse and Trento, see also http://arxiv.org/abs/1301.7741) Daniel Kressner, EPF Lausanne Title: Nonlinear eigenvalue problems with specified eigenvalues Abstract: This talk considers eigenvalue problems that are nonlinear in the eigenvalue parameter: $T(\lambda)x = 0$ for some matrix-valued function T. We are concerned with the task of assessing the quality of eigenvalue approximations obtained from a numerical algorithm, for example a quasi-Newton method. While the usual resolvent norm addresses this question for a single eigenvalue of multiplicity one, the general setting involving the simultaneous approximation of several eigenvalues is significantly more difficult. We show how the recently introduced concept of invariant pairs can be put to good use and derive a singular value optimization characterization. (Joint work with Michael Karow of TU Berlin and Emre Mengi of Koc University.) Adrian S. Lewis, Cornell University Title: Active sets and nonsmooth geometry Abstract: The active constraints of a nonlinear program typically define a surface central to understanding both theory and algorithms. The standard optimality conditions rely on this surface; they hold generically, and then the surface consists locally of all solutions to nearby problems. Furthermore, standard algorithms "identify" the surface: iterates eventually remain there. A blend of variational and semi-algebraic analysis gives a more intrinsic and geometric view of these phenomena, attractive for less classical optimization models. A recent proximal algorithm for composite optimization gives an illustration. (Joint work with J. Bolte, A. Daniilidis, D. Drusvyatskiy, M.L. Overton and S. Wright.) Zhaosong Lu, Simon Fraser University Title: Randomized block coordinate gradient methods for a class of nonlinear programming Abstract: In this talk we discuss randomized block coordinate gradient (RBCG) methods for minimizing the sum of two functions in which one of them is block separable. In particular, we present new iteration complexity results for these methods when applied to convex optimization problems. We also propose nonmonotone RBCG methods for solving a class of nonconvex problems with the above structure, and establish their global convergence. Finally, we present new complexity results for the accelerated RBCG method proposed by Nesterov for solving unconstrained convex optimization problems. Volker Mehrmann, TU-Berlin Title: Stability of descriptor systems Abstract: In this talk we discuss stability concepts for dynamical systems described by (linear or nonlinear) differential-algebraic equations (DAEs), we survey classical stability concepts (Lyapunov, Sacker Sell spectra), their generalizaions to DAEs, and also appropriate distance concepts to the nearest unstable system. We also discuss computational methods to compute the spectra and we present first results on the computation of the appropriate distance to the nearest unstable system. Julio Moro, Universidad Carlos III de Madrid Title: Directional perturbations in structured eigenproblems Abstract: The design and analysis of structure-preserving algorithms for structured eigenproblems has led in the last decades to a steady interest in structured eigenvalue perturbation theory, i.e., in analyzing the behavior of eigenvalues and other spectral objects (e.g., invariant subspaces, sign characteristics, ...) when a matrix or operator is subject to perturbations which belong to the same class of operators as the unperturbed one. It is well known that this behavior can be quite different from the behavior under arbitrary, non-structured perturbations. In this talk we make use of the Newton polygon, an elementary geometric construction first devised by Sir Isaac Newton, to give an overview of first order perturbation theory, i.e., of results involving the local variation of eigenvalues as expressed by their directional derivatives as a function of the perturbation. We will do this both when the perturbation is unstructured (i.e., arbitrary), and when it is structured, i.e., when it belongs to the same class of interest as the unperturbed matrix. This latter case shows up in many relevant practical situations when eigenvalues need to be pushed in certain specific directions, or must be moved as fast as possible away from a critical (or dangerous) region by a small, usually structured, perturbation. Special emphasis will be made on classes of matrices and matrix pencils with symmetries with respect to some indefinite scalar product, since these often arise in applications to Control and

Systems Theory.   Franz Rendl, University of Klagenfurt Title: A hierarchy of relaxations for max-cut and related problems based on small exact subproblems Abstract: The basic semidefinite relaxation for Max-Cut, underlying the Goemans-Williamson hyperplane rounding procedure, allows various tightenings. The simplest one includes constraints from the metric polytope. More refined approaches are iterative, and provide a sequence of relaxations, which come arbitrarily close to the convex hull of cut matrices, but at an increasingly high computational effort. A natural systematic hierarchy was introduced by Lasserre. The first step in this hierarchy corresponds to the basic semidefinite relaxation, where the optimization is done over the set of correlation matrices. The second one is a relaxation in the space of matrices of order O(n^2). We propose an iterative refinement of the basic semidefinite relaxation intersected with the metric polytope. The refinement is obtained by asking that submatrices to certain carefully selected small node-induced subgraphs of the problem are actually contained in the cut-polytope of the subgraph. The matrix order in the refinement process is always n, and only the number of constraints may grow exponentially. A similar approach has been used in the context of linear relaxations under the heading of 'target cuts'. We will discuss the connections and differences of the present approach to previous methods, and also consider applications to stable-set relaxations. We provide some theoretical insights as well as first computational experience. (Joint work with E. Adams, M. Anjos of Montreal and A. Wiegele of Klagenfurt.)   Michael Saunders, Stanford University Title: Solving linear systems by orthogonal tridiagonalization (GMINRES and/or GLSQR) Abstract: A general matrix A can be reduced to tridiagonal form by orthogonal transformations on the left and right: U^T AV = T. We can arrange that the first columns of U and V are proportional to given vectors b and c. An iterative form of this process was given by Saunders, Simon, and Yip (SINUM 1988) and used to solve square systems Ax = b and A^T y = c simultaneously. (One of the resulting solvers becomes MINRES when A is symmetric and b = c.) The approach was rediscovered by Reichel and Ye (NLAA 2008) with emphasis on rectangular A and least-squares problems Ax \approx b. The resulting solver was regarded as a generalization of LSQR (although it doesn't become LSQR in any special case). Careful choice of c was shown to improve convergence. In his last year of life, Gene Golub became interested in ``GLSQR'' for estimating c^T\! x and b^T\! y without computing x or y. Golub, Stoll, and Wathen (ETNA 2008) revealed that the orthogonal tridiagonalization is equivalent to a certain block Lanczos process. This reminds us of Golub, Luk, and Overton (TOMS 1981): a block Lanczos approach to computing singular vectors.   Josef Sifuentes, Texas A&M University Title: Approximate Murphy-Golub-Wathen preconditioning for saddle point problems Abstract: Murphy, Golub, and Wathen proposed a preconditioner for saddle-point systems that yields a diagonalizable coefficient matrix having three distinct eigenvalues, giving exact convergence of GMRES in three iterations. However, this preconditioner involves the inverse of a large submatrix. Practical computations only approximate this inverse, so GMRES will generally require more iterations. How many more? Recent results on the stability of GMRES lead to rigorous bounds on the number of required iterations as a function of the accuracy to which the preconditioner is applied, along with spectral properties of the constituent matrices. Numerical computations verify these results for problems from optimization and fluid dynamics.   Daniel B. Szyld, Temple University Title: Inexact and truncated parareal-in-time Krylov subspace methods for parabolic optimal control problems Abstract: We study the use of inexact and truncated Krylov subspace methods for the solution of the linear systems arising in the discretized solution of the optimal control of a parabolic partial differential equation. An all-at-once temporal discretization and a reduction approach are used to obtain a symmetric positive definite system for the control variables only, where a Conjugate Gradient (CG) method can be used at the cost of the solution of two very large linear systems in each iteration. We propose to use inexact Krylov subspace methods, in which the solution of the two large linear systems are not solved exactly, and their approximate solutions can be progressively less exact. The option we propose is the use of the parareal-in-time algorithm for approximating the solution of these two linear systems. The use of less parareal iterations makes it possible to reduce the time integration costs and to improve the time parallel scalability, and therefore, making it possible to really consider optimization in real time. We also show that truncated methods could be used without much delay in convergence, but with important savings in storage. Spectral bounds are provided and numerical experiments with the full orthogonalization method (FOM), and with inexact and truncated version of FOM are presented, illustrating the potential of the proposed methods. Lloyd. N. Trefethen, University of Oxford Title: Eigenvalue avoidance Abstract: The indefinite least squares problem and the equality constrained indefinite least squares problem are modifications of the least squares problem and the equality constrained least squares problem, respectively, involving the minimization of a certain type of indefinite quadratic form. Such problems arise in the solution of Total Least Squares problems, in parameter estimation and in H-infinity smoothing. Algorithms for computing the numerical solution of indefinite least squares and indefinite least squares with equality constraint are described by Bojanczyk et al. and Chandrasekharan et al. The indefinite least squares problem and the equality constrained indefinite least squares problem can be expressed in an equivalent fashion as augmented square linear systems. Exploiting the particular structures of the coefficient matrices of such systems, new algorithms for computing the solution of such problems are proposed relying on the anti-triangular factorization of the coefficient matrix (by Mastronardi et al.). Some results on their stability are shown together with some numerical examples. (Joint work with Nicola Mastronardi.) Charles Van Loan, Cornell University Title: The higher order generalized singular value decompositions Abstract: Suppose you have a collection of data matrices each of which has the same number of columns. The HO-GSVD can be used to identify common features that are implicit across the collection. It works by identifying a certain (approximate) invariant subspace of a matrix that is a challenging combination of the collection matrices. In describing the computational process I will talk about the Higher Order CS decomposition and a really weird optimization problem that I bet you have never seen before! (Joint work with Orly Alter, Priya Ponnapalli, and Mike Saunders.) Stephen Vavasis, University of Waterloo Title: Convex relaxation for finding planted influential nodes in a social network Abstract: We consider the problem of maximizing influence in a social network. We focus on the case that the social network is a directed bipartite graph whose arcs join senders to receivers. We consider both the case of deterministic networks and probabilistic graphical models, that is, the so-called ``cascade'' model. The problem is to find the set of the k most influential senders for a given integer k. Although this problem is NP-hard, there is a polynomial-time approximation algorithm due to Kempe, Kleinberg and Tardos. In this work we consider convex relaxation for the problem. We prove that convex optimization can recover the exact optimizer in the case that the network is constructed according to a generative model in which influential nodes are planted but then obscured with noise. We also demonstrate computationally that the convex relaxation can succeed on a more realistic generative model called the â€œforest fireâ€• model. (Joint work with Lisa Elkin and Ting Kei Pong of Waterloo.) Henry Wolkowicz, University of Wtareloo Title: Taking advantage of degeneracy in cone optimization with applications to sensor network localization and molecular conformation Abstract: The elegant theoretical results for strong duality and strict complementarity for linear programming, LP, lie behind the success of current algorithms. However, the theory and preprocessing techniques that are successful for LP can fail for cone programming over nonpolyhedral cones. Surprisingly, many instances of semidefinite programming, SDP, problems that arise from relaxations of hard combinatorial problems are degenerate. (Slater's constraint qualification fails.) Rather than being a disadvantage, we show that this degeneracy can be exploited. In particular, several huge instances of SDP completion problems can be solved quickly and to extremely high accuracy. In particular, we illustrate this on the sensor network localization and Molecular conformation problems. Margaret Wright, New York University Title: Numerical linear algebra and derivative-free optimization on Facebook: "in a relationship" or "it's complicated"? Abstract: The longstanding connection between numerical linear algebra and optimization has often, but not always, been close---think about linear programming, the product form of the inverse, and numerically stable rank-one updates of the basis. If we pretend that derivative-free optimization and numerical linear algebra are on Facebook (after all, SIAM is), how would they describe their connection? This talk will discuss a selection of the varied roles played by numerical linear algebra in recent work on derivative-free optimization. Jane Ye, University of Victoria Title: Minimizing the condition number to construct optimal experimental designs Abstract: Many practical and theoretical problems in science and engineering consider the relationship between an explanatory variable and a response variable as a p-th order polynomial regression model. In a controlled experiment, one can select the levels of the explanatory variable and the way to select the level of the explanatory variables is called a design of experiment. In the literature, several design criteria have been introduced for constructing optimal experimental designs so that the least squares estimator is more efficient. We introduced a new optimality criterion, the K-optimality criterion, for constructing optimal experimental designs for polynomial regression models based on minimizing the condition number of the information matrix. A condition number of a symmetric positive definite matrix is the ratio of the largest and the smallest eigenvalues. It is an important measure of sensitivity for many practical problems. It is well known that the condition number for a positive definite matrix as the ratio of the maximum eigenvalue to the minimum eigenvalue is usually nonsmooth. For the case where the design space is the interval [-1,1], we show that the condition number of the information matrix is continuously differentiable. Moreover we show that there is always a symmetric K-optimal design with exactly p+1 support points including the boundary points -1 and 1. These results allow us to use standard optimization solvers to find the K-optimal designs. For the general case where the design space is an interval [a,b], we propose a two-stage strategy. In the first stage we find an optimal moment matrix by using semi-definite program (SDP). In the second stage we use the entropy optimization to find an approximate optimal design.

Links:
Workshop webpage, http://www.pims.math.ca/scientific-event/130808-wnlao Slides for some of the talks available at
http://www.mathtube.org/conference/workshop-numerical-linear-algebra-and-optimization

---

File Uploads:
Additional Upload 1: http://www.pims.math.ca/files/final_report/greenbaum_abstract.pdf

---