Acceleration and Extrapolation Methods
Poster Session Abstracts
Tuesday, July 25, 2023

Newton-Anderson at Singular Points
Matt Dallas

We develop convergence and acceleration theory for Anderson acceleration applied to Newton's method for nonlinear systems in which the Jacobian is singular at a solution. For these problems, the standard Newton algorithm converges linearly in a region about the solution; and, it has been previously observed that Anderson acceleration can substantially improve convergence without additional a priori knowledge, and with little additional computation cost. We present an analysis of the Newton-Anderson algorithm in this context, and introduce a novel and theoretically supported safeguarding strategy. The convergence results are demonstrated with several benchmark examples, including the Chandrasekhar H-equation.

Stopping Criteria for the Conjugate Gradient Algorithm in High-order Finite Element Methods
Yichen Guo

Solving partial differential equations (PDE) involves discretizating the PDE and solving the linear system. Determining when to stop the iteration of the linear solver is essential to the computational efficiency. A commonly used stopping criterion suggests stopping when the relative residual norm is less than a tolerance, which often leads to over-solving. We present new stopping criteria that balance algebraic and discretization errors for the conjugate gradient algorithm applied to high-order finite element discretizations of Poisson problems. Numerical experiments, including tests with anisotropic meshes and highly variable piecewise constant coefficients, demonstrate that the proposed criteria efficiently avoid both premature termination and over-solving.

A Unified Framework of Anderson Acceleration for Operator Splitting
Qiang Heng

Anderson acceleration is a powerful technique for accelerating the convergence of fixed-point iterations leveraging past iterates. In this paper, we present a generalization of the Anderson accelerated Douglas-Rachford splitting algorithm proposed by Fu, Zhang, and Boyd. Our work introduces a unified framework of type-II Anderson acceleration that encompasses multiple prevalent operator splitting methods. This generalization enhances the versatility and applicability of Anderson acceleration in practice. By considering mild assumptions required by the original splitting methods, we establish global asymptotic convergence to an optimal point. We conduct a diverse set of numerical experiments to showcase the algorithm's ability to achieve both global convergence and notable speed improvements, making it a promising tool for accelerating the convergence of fixed-point iterations in modern convex optimization.

A Mini-Batch Quasi-Newton Proximal Method for Constrained Total-Variation Nonlinear Image Reconstruction
Tao Hong

Over the years, computational imaging with accurate nonlinear physical models has drawn considerable interest due to its ability to achieve high-quality reconstructions. However, such nonlinear models are computationally demanding. A popular choice for solving the corresponding inverse problems is accelerated stochastic proximal methods (ASPMs), with the caveat that each iteration is expensive. To overcome this issue, we propose a mini-batch quasi-Newton proximal method (BQNPM) tailored to image-reconstruction problems with total-variation regularization. It involves an efficient approach that computes a weighted proximal mapping at a cost similar to that of the proximal mapping in ASPMs. However, BQNPM requires fewer iterations than ASPMs to converge. We assess the performance of BQNPM on three-dimensional inverse-scattering problems with linear and nonlinear physical models. Our results on simulated and real data show the effectiveness and efficiency of BQNPM.

Vector RRE for Classifying Chaos in Symplectic Maps
Maximilian Ruth

Many important qualities of nuclear plasma confinement devices can be determined via the Poincare plot of a symplectic return map. These qualities include the locations of the magnetic core (good for nuclear fusion), chaotic regions (bad for confinement), and magnetic islands (potentially good for ejection of particles). The convergence rate of ergodic averages have been shown to successfully categorize these orbits, but many iterations of the return map are needed to implement this naively. Recently, it has been shown that a weighted average can be used to accelerate the convergence, resulting in a useful method for categorizing trajectories.

In this poster, we will show how the vector reduced rank extrapolation method (vector RRE) can classify chaos using fewer iterations of the map than the non-adaptive weighted average methods. Additionally, when the trajectory is not chaotic, we will show that the eigenmodes of the learned extrapolation model can be used to recover the geometry of the Poincare plot. This includes determining both the number of islands in a chain and the rotation number, an important quantity in both plasma physics and dynamical systems. A theorem for the convergence of vector RRE for this problem is presented, where the proof relies on techniques from both Diophantine number approximation and numerical linear algebra.

Simple Extrapolation for Tensor Eigenvalue Problems
Rhea Shroff

Tensors are algebraic objects that define a multi-linear relationship between sets of algebraic objects related to a vector space. Hence, they naturally arise in data-intensive problems. With the advent of increasingly complex data sets, there is a need to make reasonable inferences from them. Due to the scale of the problem, computation can easily be limited by the time it takes to conduct operations on it. This is why we require extrapolation methods to improve performance. Extrapolation methods are an extremely powerful tool used to increase speed and accuracy for various numerical methods. For this poster, I will focus on a simple extrapolation method for the shifted symmetric higher order power method for tensor eigenvalue problems.

An adaptive superfast inexact proximal augmented Lagrangian method for smooth nonconvex composite optimization problems

Arnesh Sujanani

This work presents an adaptive superfast proximal augmented Lagrangian (AS-PAL) method for solving linearly-constrained smooth nonconvex composite optimization problems. Each iteration of AS-PAL inexactly solves a possibly nonconvex proximal augmented Lagrangian (AL) subproblem obtained by an aggressive/adaptive choice of prox stepsize with the aim of substantially improving its computational performance followed by a full Lagrangian multiplier update. A major advantage of AS-PAL compared to other AL methods is that it requires no knowledge of parameters (e.g., size of constraint matrix, objective function curvatures, etc) associated with the optimization problem, due to its adaptive nature not only in choosing the prox stepsize but also in using a crucial adaptive accelerated composite gradient variant to solve the proximal AL subproblems. The speed and efficiency of AS-PAL is demonstrated through extensive computational experiments showing that it can solve many instances more than ten times faster than other state-of-the-art penalty and AL methods, particularly when high accuracy is required.

Numerical Study of Anderson Acceleration

Ning Wan

Anderson Acceleration is a method which allows to speed-up the convergence of the fixed point iteration, which is widely used within the SCF (Self-Consistent Field Iteration) method in quantum chemistry/physics computations. In this poster presentation, we aim to explore Anderson Acceleration under perturbations and various truncation approaches. By studying the effects of perturbations and different truncation methods, we can gain insights into the behavior and performance of Anderson Acceleration in practical applications.