Asymptotically improved solvers for the variable coefficient Helmholtz equation

James Bremer
Department of Mathematics
University of Toronto

Existing numerical methods for the variable coefficient Helmholtz equation are prohibitively expensive in the high-frequency regime — that is, when the wavenumber k of the problem is large. A fundamental obstacle is that in d dimensions, standard discretization methods require $\mathcal{O}(k^d)$ points to represent it solutions. The computational cost to solve the Helmholtz equation must then be at least $\mathcal{O}(k^d)$ and is, in fact, often far larger due to the difficulties which arise when solving the linear systems which arise in the highly oscillatory regime. We will discuss two cases in which asymptotically improved solvers can be constructed. First, we will describe a solver for the one-dimensional variable coefficient Helmholtz equation which runs in time independent of the wavenumber. Then, we will describe a method for simulating scattering from a two dimensional radially symmetric potential which runs in $\mathcal{O}(k \log(k))$ time.

Aaron Baier-Reinio
ambaierr@uwaterloo.ca
University of Waterloo
Department of Applied Mathematics

Error analysis of divergence-free discontinuous Galerkin methods for incompressible flow problems under minimal regularity

The standard a priori error analysis of discontinuous Galerkin (DG) methods for the Stokes problem requires that the exact velocity field $\boldsymbol{u}$ is sufficiently regular, i.e. $\boldsymbol{u} \in \boldsymbol{H}^2(\Omega)$. In this work, we give an a priori error analysis of two lowest-order hybridizable DG methods for the Stokes problem, while making only minimal regularity assumptions on the exact solution. In particular, our analysis requires only the fractional Sobolev regularity $\boldsymbol{u} \in \boldsymbol{H}^{1+s}(\Omega)$ for any real number $s \in [0, 1]$. Moreover, the numerical methods under consideration have previously been shown to produce $\boldsymbol{H}(\mathrm{div})$-conforming and divergence-free approximate velocity fields, and we use these properties to obtain optimal velocity error estimates that are independent of the pressure. Finally, we discuss how this minimal regularity error analysis can be extended to divergence-free DG methods for the nonlinear steady Navier–Stokes problem.

Joint work with Sander Rhebergen, University of Waterloo and Garth N. Wells, University of Cambridge.

Strong imposition of the no-slip boundary condition for the compressible Navier-Stokes equations

Anita Gjesteland[1] and Magnus Svärd[1]

[1]University of Bergen

We consider herein the compressible Navier-Stokes equations augmented with no-slip adiabatic wall boundary conditions. We propose a numerical scheme to approximate the equations, based on the second-order summation-by-parts (SBP) operator. Our focus is the stability properties of the scheme, with special emphasis on the chosen boundary treatment. In particular, the no-slip condition is imposed strongly by injection while the temperature condition is imposed weakly using the simultaneous approximation term (SAT) technique. By first linearising the scheme, we prove that this boundary procedure leads to a linearly (energy) stable scheme. Next, by assuming that the interior scheme is entropy stable, we show that the chosen combination of boundary impositions also yields a non-linearly (entropy) stable scheme. Some numerical simulations are included to demonstrate the robustness of the boundary treatment.

DGMT: A semidiscretization method for solving high-dimensional parabolic partial differential equations with deep learning

Short title: Deep Galerkin Method with Timestepping

Ruining Wu[a] and Christina Christara[b]
[a]University of Toronto, Department of Computer Science, rwu@cs.toronto.edu
[b]University of Toronto, Department of Computer Science, ccc@cs.toronto.edu

Partial Differential Equations (PDEs) are ubiquitous tools for modelling systems in the physical world, and high-dimensional PDEs naturally occur when the dimensionalities of the systems are large. One example is the pricing of multi-asset options using the Black-Scholes model in finance, where each underlying asset corresponds to a spatial dimension of the PDE. Solving high-dimensional PDEs with standard numerical methods such as Finite Differences or Finite Elements lead to the curse of dimensionality; that is, the number of unknowns scales exponentially with the number of dimensions, making computation infeasible. Recently, neural network methods such as [Sirignano and Spiliopoulos, 2018]'s Deep Galerkin Method (DGM) have overcome the curse of dimensionality, and thus make solving high-dimensional PDEs computationally feasible. We introduce the Deep Galerkin Method with Timestepping (DGMT), which improves upon the DGM by applying a time-discretization to the PDE of interest and solving the spatial dimensions with a neural network. In the DGMT, solution values at a timestep are directly dependent upon solution values at a previous timestep. Numerical results show that the DGMT produces a more accurate pointwise result for similar computational work compared to the DGM. We also show that DGMT has other desirable properties such as reduced fluctuations in the predicted value every training epoch. Finally, the DGMT allows error to be measured as a function of the stepsize, which allows us to study convergence of the method in the traditional sense of numerical methods.

Efficient Pricing and Hedging of High Dimensional American Options using Deep Recurrent Networks

Andrew S. Na[a] and Justin W.L. Wan[b]

[a]University of Waterloo, Cheriton School of Computer Science; [b]University of Waterloo, Cheriton School of Computer Science

ARTICLE HISTORY
Compiled May 11, 2022

Abstract
We propose a deep Recurrent neural network (RNN) framework for approximating high dimensional American options prices and deltas. The architecture of the framework uses two deep RNNs, where one network learns the price and the other learns the delta of the option for each timestep. Our proposed network computes prices and deltas over the entire spacetime, not only at a given point (e.g. $t = 0$). The computational cost of our proposed approach is linear in time, which improves on the quadratic time seen for feedforward networks that price American options. The computational memory cost of our method is constant in memory, which is an improvement over the linear memory costs seen in feedforward networks. We demonstrate our contributions using numerical simulations, and show that the proposed deep RNN framework is computationally more efficient than traditional feedforward neural network frameworks in time and memory.

CONTACT Andrew S. Na. Email: andrew.na@uwaterloo.ca

A high-order deferred correction method for the solution of free boundary problems using penalty iteration, with an application to American option pricing

Short title: A High Order Method for Solving Free Boundary Problems

Dawei Wang[a], Kirill Serkh[b] and Christina Christara[c]
[a]University of Toronto, Department of Computer Science, dwang@cs.toronto.edu
[b]University of Toronto, Department of Mathematics, kserkh@math.toronto.edu
[c]University of Toronto, Department of Computer Science, ccc@cs.toronto.edu

This paper presents a high-order deferred correction algorithm combined with penalty iteration for solving free and moving boundary problems, using a fourth-order finite difference method. Typically, when free boundary problems are solved on a fixed computational grid, the order of the solution is low due to the discontinuity in the solution at the free boundary, even if a high-order method is used. Using a detailed error analysis, we observe that the order of convergence of the solution can be increased to fourth-order by solving successively corrected finite difference systems, where the corrections are derived from the previously computed lower order solutions. The penalty iterations converge quickly given a good initial guess. We demonstrate the accuracy and efficiency of our algorithm using several examples. Numerical results show that our algorithm gives fourth-order convergence for both the solution and the free boundary location. We also test our algorithm on the challenging American put option pricing problem. Our algorithm gives the expected high-order convergence.

Optimal Asset Allocation For Outperforming A Stochastic Benchmark Target

Chendi Ni[1], Yuying Li[2], Peter Forsyth[3], and Ray Carroll[4]

[1]Cheriton School of Computer Science, University of Waterloo, Waterloo, N2L 3G1, Canada, chendi.ni@uwaterloo.ca
[2]Cheriton School of Computer Science, University of Waterloo, Waterloo, N2L 3G1, Canada, yuying@uwaterloo.ca
[3]Cheriton School of Computer Science, University of Waterloo, Waterloo, N2L 3G1, Canada, paforsyt@uwaterloo.ca
[4]Neuberger Berman Breton Hill, Toronto, M4W 1A8, Canada, Ray.Carroll@nb.com

April 27, 2022

Abstract

We propose a data-driven Neural Network (NN) optimization framework to determine the optimal multi-period dynamic asset allocation strategy for outperforming a general stochastic target. We formulate the problem as an optimal stochastic control with an asymmetric, distribution shaping, objective function. The proposed framework is illustrated with the asset allocation problem in the accumulation phase of a defined contribution pension plan, with the goal of achieving a higher terminal wealth than a stochastic benchmark. We demonstrate that the data-driven approach is capable of learning an adaptive asset allocation strategy directly from historical market returns, without assuming any parametric model of the financial market dynamics. The optimal adaptive strategy outperforms the benchmark constant proportion strategy, achieving a higher terminal wealth with a 90% probability, a 46% higher median terminal wealth, and a significantly more right-skewed terminal wealth distribution.

Keywords Dynamic Asset Allocation, Data Driven, Neural Networks

Accelerating Your Research Through Advanced Research Computing

John Morton
Director of Technology
SHARCNET

We will present the resources available to Canadian researchers through the national Advanced Research Computing (ARC) platform. Resources include hardware such as CPUs, GPUs, storage but also software, services and training and user support. All aspects of the national platform are available at no cost to you.

Simulating the aggregation of micro-organisms

Lennaert van Veen
Ontario Tech University

Collective behaviour of organisms is readily observed in nature, from the murmuration of starlings to ants forming supply lines. How such organization emerges from the rules that individual organisms follow and the way they interact is a long-standing question. There is no overarching theoretical description of emergent behaviour across space and time scales and organisms. Most research is either based on experimental observations or on computer simulations. These simulations come in two forms: agent based and continuous. Agent-based modelling has the advantage that rules of communications and propulsion can be accurately prescribed. A downside is that extreme computational resources are often required to simulate groups of organisms large enough to exhibit collective behaviour. Continuous models describe the evolution of the density of organisms and thus constitute a kind of "thermodynamic limit" of agent-based models. We consider a continuous model in two spatial dimensions originally formulated by Razvan Fetecau and tune its attraction, repulsion and alignment interactions to mimic those of aggregating micro-organisms. Furthermore, we introduce a strongly nonlinear saturation effect that is know to lead to a rich dynamical repertoire in one spatial dimension. The goal is to simulate the formation of clusters and study the distribution of cluster sizes, their stability and other questions currently out of reach of agent-based simulations. Since the model takes the form of a non-local integro-differential equation, its numerical treatment is not straightforward and relies on convolutions theorems and exponential time differencing. Joint work with Andree Qi (MSc student at Ontario Tech) and Luciano Buono (UQAR).

MM-ADMM: Implicit Integration of MMPDEs in Parallel

Short title: MM-ADMM

Connor Tannahill[a] and Justin Wan[b]
[a]University of Waterloo, David R. Cheriton School of Computer Science, ctannahi@uwaterloo.ca
[a]University of Waterloo, David R. Cheriton School of Computer Science, justin.wan@uwaterloo.ca

Variational mesh adaptation seeks to generate meshes which are well adapted to the underlying computational problem. In particular, methods based on the moving mesh PDE (MMPDE) methods are a class of computational methods which generate the optimal meshes as the minimizer to meshing functionals which describe the desired properties that the mesh should satisfy. Recent work in this area has improved the computational efficacy of these methods, with one focus in particular being the development of parallel methods. We introduce a novel optimization-based approach for variational mesh adaptation based on recent practical MMPDE methods combined with recent techniques for quickly solving large-scale, non-linear constraint problems with scalable parallelization. We call this method MM-ADMM. The resulting method bears resemblance to meshing algorithms based on the spring-analogy, while producing high-quality adaptive meshes using more principled variational models. We demonstrate the advantages of our method over standard MMPDE methods for generating two and three dimensional meshes parallel.

Variational Data Assimilation

Hamidreza Moazzami, Nicholas Kevlahan

McMaster University, 1280 Main St West, Hamilton, ON, Canada

Data assimilation (DA) improves numerical simulations of physical phenomena such as weather forecasting by optimally incorporating observations. We present results obtained by applying four-dimensional variational data assimilation (4D-VAR) to the one-dimensional linear advection equation discretized using the Lax–Wendroff scheme. We compare the performance of two versions of 4D-VAR, where the discrete adjoint model is: 1. derived directly from the discrete equations, or 2. obtained by discretizing the continuous form of the adjoint equations. We find that the two approaches give equivalent results. We also observed that assuming that the error covariance matrix is the identity matrix is sufficient for effective error reduction. However, if the number of observations is too low (e.g. less than 3 for the grid of size 32), there is insufficient information about the background state which degrades the accuracy of the DA.

Energy-conserved FDTD methods with local mesh refinement for Maxwell's equations with Drude model

Cong Wang

Department of Mathematics and Statistics
York University

Maxwell's equations with Drude model describe the electromagnetic wave propagation in metamaterials. When the domain of interest contains local solution behavior, the computation of problem is required to improve the resolution capacity and efficiency of solution. However, in the presence of local mesh refinement, it is important and difficult to provide the suitable interface scheme/condition which can preserve energy conservation and guarantee high accuracy. In this talk, I will first present our development of energy-preserving local-refined FDTD schemes for solving Maxwell's equations with Drude model. Two second-order space accurate schemes based on energy-conserved splitting FDTD schemes are constructed. The interface equations are carefully treated so that two algorithms are energy conservative, stable and accurate for long time simulation. Fast implementations of solving the schemes are also considered. Numerical experiments are then reported. This is joint work with Prof. Dong Liang.

MechFacility3D: Simulating Multibody Dynamics in 3D

Dylan Bassi
McMaster University, Computing and Software

Abstract: We are developing MechFacility3D, a C++ package for the modelling and simulation of multibody systems (MBS) in Cartesian 3D coordinates. Starting from a human-readable mechanism specification in YAML format, we construct a Lagrangian function and constraints. Through automatic differention, the equations of motion are constructed "behind-the-scenes", which usually results in an index-3 differential-algebraic equation (DAE). We numerically integrate the resulting DAE by the high-index DAE solver DAETS of Nedialkov and Pryce (NP). Visualization occurs through NP's `animatemech` MATLAB program.

The goal of this project is to extend the functionality of NP's 2D Mechanism Facility to incorporate 3D, as well to introduce new constraints and dissipative elements. We present a theoretical background, overview of our system as well as demonstrate its functionality through examples.

HARFE: Hard-Ridge Random Feature Expansion

Esha Saha[1], Hayden Schaeffer[2], and Giang Tran[1]

[1]Department of Applied Mathematics, University of Waterloo
[2]Department of Mathematical Sciences, Carnegie Mellon University

Abstract

We propose a random feature model for approximating high-dimensional sparse additive functions called the hard-ridge random feature expansion method (HARFE). This method utilizes a hard-thresholding pursuit based algorithm applied to the sparse ridge regression (SRR) problem to approximate the coefficients with respect to the random feature matrix. The SRR formulation balances between obtaining sparse models that use fewer terms in their representation and ridge-based smoothing that tend to be robust to noise and outliers. In addition, we use a random sparse connectivity pattern in the random feature matrix to match the additive function assumption. We prove that the HARFE method is guaranteed to converge with a given error bound depending on the noise and the parameters of the sparse ridge regression model. Based on numerical results on synthetic data as well as on real datasets, the HARFE approach obtains lower (or comparable) error than other state-of-the-art algorithms.