

SONAD 2026: 41st Southern Ontario Numerical Analysis Day

Friday, May 1, 2026
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada

Speakers, titles and abstracts of talks¹

Ricardo Baptista

Statistical Sciences, University of Toronto
r.baptista@utoronto.ca

Dynamics and Memorization Behaviour of Generative Diffusion Models

Diffusion models have emerged as a powerful framework for generative modeling in the information sciences and many scientific domains. To generate samples from the target distribution, these models rely on learning the the gradient of the data distribution's log-density using a score matching procedure. A key element for the success of diffusion models is that the optimal score function is not identified when solving the denoising score matching problem. In fact, the optimal score in both unconditioned and conditioned settings leads to a diffusion model that returns to the training samples and effectively memorizes the data distribution. In this presentation, we study the dynamical system associated with the optimal score and describe its long-term behavior relative to the training samples. Lastly, we show the effect of two forms of score function regularization on avoiding memorization: restricting the score's approximation space and early stopping of the training process. These results are numerically validated using distributions with and without densities including image-based inverse problems for scientific machine learning applications.

Hanno Rein

Astronomy and Astrophysics / Computer Science, University of Toronto
hanno.rein@utoronto.ca

Orbits, Chaos, Instabilities, Machine Learning and how to write the fastest N-body code in the world

I will start with a short introduction to orbital mechanics, the N-body problem, and the kinds of algorithms we can use to numerically solve the N-body problem. I will then discuss the implementation of WHFast512, an N-body code developed by my research group. It is written in assembly and uses AVX512 SIMD instructions. I will present some interesting results from this field, including the discovery that there is a finite probability that the Earth will collide with another planet in the next five billion years. I will finish by touching on the question of whether machine learning can help us determine orbital parameters of newly discovered extrasolar planets.

¹All talks take place in the Bahen Centre for Information Technology Building, **room BA 1190**, 40 St. George St, Toronto, Ontario, which is located in the University of Toronto Saint George campus.

Xinwen Ding

Ph.D. student (with Adam Stinchcombe)

Mathematics, University of Toronto

xinwen.ding@mail.utoronto.ca

Walk-on-Interfaces: A Monte Carlo Estimator for Elliptic Interface Problem

Elliptic interface problems arise in many areas of science and engineering, modeling heterogeneous materials whose physical properties change abruptly across internal boundaries. Computing solutions to these problems efficiently and accurately remains challenging, especially in domains with multiple irregular interfaces. In this poster, we present Walk-on-Interfaces (WoI), a grid-free Monte Carlo estimator for Neumann elliptic interface problems with general flux jump conditions. Unlike many numerical schemes, WoI maintains uniform accuracy throughout the domain and avoids near-interface singularities. Moreover, gradients of the solution can be estimated at almost no additional cost by differentiating the Green's function within WoI. Taking a scientific machine learning approach, we train a deep neural network to filter out high-frequency sampling noise, yielding a smooth and continuous representation of the solution. The resulting method is highly parallelizable, scales naturally to high dimensions, and can solve problems that are intractable for traditional numerical solvers. Numerical experiments demonstrate the effectiveness of the approach and highlight its potential for real-world applications.

Fabian Bleitner

Postdoctoral fellow (with Bartosz Protas)

Mathematics & Statistics, McMaster University

bleitnef@mcmaster.ca

Probing the sharpness of bounds related to the Ladyzhenskaya-Prodi-Serrin conditions

In this talk we examine the sharpness of regularity estimates associated with the Ladyzhenskaya-Prodi-Serrin conditions for the 3D Navier-Stokes equations. These conditions assert that the solution u remains smooth on $[0, T]$ if

$$\int_0^T \|u\|_{L^q}^p dt < \infty$$

for $\frac{2}{p} + \frac{3}{q} \leq 1$ and $q > 3$. At the same time, the rate of growth of the Lebesgue norm is subject to the bound

$$\frac{d}{dt} \|u\|_{L^q} \leq c \|u\|_{L^q}^{\frac{3q-1}{q-3}}.$$

If this a priori bound is attained on $[0, T]$, the integral criterion fails, implying blow-up. To probe whether this bound is sharp, we formulate an optimization problem where one maximizes the growth rate of the L^q -norm subject to suitable constraints. This problem is solved numerically using a Riemannian conjugate gradient approach, producing candidates for initial data that may lead to singularity formation in the Navier-Stokes equations. Preliminary computations indicate saturation of the bound for certain values of q .

Ray Wu

Ph.D. student (with Christina Christara)
Computer Science, University of Toronto
rwu@cs.toronto.edu

Quantization error under high-order convolution smoothing in parabolic PDE models for finance

Numerical solution of parabolic PDEs with nonsmooth initial conditions often exhibits reduced and/or unstable convergence orders due to the so-called quantization error. For a class of nonsmooth initial conditions, including representative payoffs of financial derivatives, we study how high-order convolution-based smoothings modify the quantization error. For each initial condition and smoothing operator, we derive closed-form expressions for the leading order terms of the quantization error. Examining the dependence of the quantization error leading coefficient on the relative location of the discontinuity with respect to the discretization grid, we identify the minimum smoothing order required to guarantee stable fourth-order convergence (with or without grid alignment), when using fourth-order discretizations in space and time. These results explain when extrapolation-based acceleration methods, including Richardson extrapolation and the sparse grid combination approach, succeed or fail.

Yifan Qu

Ph.D. student (with Christina Christara)
Computer Science, University of Toronto
quyf@cs.toronto.edu

Numerical Approximation of Exercise Boundary of American Options

While there are several numerical methods that have been developed to compute the value of an American option, as well as studies on the accuracy of these methods, comparatively less attention has been devoted to accurately calculating the associated optimal exercise boundary. In practice, simple techniques often produce a piecewise or “stepwise” approximation of the free boundary, which may lead to non-monotonic behavior and oscillations or lack of smoothness over time. We investigate several techniques for approximating the exercise boundary based on different interpolation strategies applied to the option value or derivative values within the region of the partial differential equation. These interpolants are then extrapolated towards the exercise region, and the free boundary is approximated by enforcing either continuity or continuity of derivative conditions. We seek a method that can attain a reasonably stable and reliable order of convergence and monotonicity with respect to time.

Joint work with Christina Christara.

Lucas Dutton

Ph.D. student (with Christopher Anand)
Computing and Software, McMaster University
duttonl@mcmaster.ca

Using Fixed-Point Approximations to Accelerate Integer Division

Most cryptographic computations are integer computations, but we will show an example where approximations of real-number computations are an efficient method of computation. Specifically, we will show that integer division can be efficiently computed using fixed-point approximations to real numbers, with a careful error analysis. The resulting implementation using RISC-V vector instructions is both much more efficient than pure-integer algorithms, and is constant-time, which is required to avoid timing attacks.

Thanh Huynh

Masters student (with Lia Bronsard and Dominik Stantejsky)
Mathematics, McMaster University
huynht17@mcmaster.ca

A numerical approach for local isoperimetric partitions

We present a numerical method for isoperimetric partitioning problems with volume constraints. The method is based on a total variation formulation to measure the perimeter of the sets, leading to a non-smooth, nonconvex optimization problem with multiple constraints. We solve the problem using an ADMM-based algorithm, with FFT techniques used to accelerate the computations. We validate the approach against known solutions, including the disk and lens cluster, and investigate its convergence behaviour numerically. Finally, we propose several conjectures about minimizing configurations in cases where the optimal shapes are unknown.

Joint work with Lia Bronsard and Dominik Stantejsky.

Shukai Chen

Ph.D. student (with Kirill Serkh)
Computer Science, University of Toronto
shukai.chen@mail.utoronto.ca

The univariate adaptive Levin method

The Levin method is a well-known technique for evaluating univariate oscillatory integrals, which operates by solving the so-called Levin ODE in order to construct an antiderivative of the integrand. It was long believed that this approach suffers from "low-frequency breakdown," meaning that the accuracy of the calculated value of the integral deteriorates when the integrand is only slowly oscillating. We recently proved that, under mild conditions that guarantee the existence of a slowly-varying solution to the Levin ODE, one can recover a potentially different slowly-varying solution to the Levin ODE if a Chebyshev spectral method is used to discretize the ODE and the resulting linear system is solved via a truncated singular value decomposition. We show that, if such univariate Levin method is combined with an adaptive subdivision strategy, it can be used to efficiently evaluate univariate oscillatory integrals regardless of the behavior of the integrand. We present the results of numerical experiments which illustrate the

consequences of our analysis and demonstrate the properties of the adaptive Levin method.

Samira Jamalidinan

Ph.D. student (with Kazem Cheshmi)

Electrical and Computer Engineering, McMaster University

jamalids@mcmaster.ca

Adaptive Retrieval and Shot Selection for Tabular Prediction

Tabular prediction is a critical task across numerous applications. The recent success of large language models has sparked various approaches for adapting them to the tabular domain. A prevalent strategy involves training or fine-tuning specialized Tabular Foundation Models (TFMs). However, TFMs require substantial computational resources, and frequent model retraining is often impractical. In-context learning (ICL), specifically, few-shot prompting, offers a resource-efficient alternative to enhance performance. Yet, identifying the most relevant rows to serve as shots remains a challenge for tabular data. This paper introduces ARASH (Adaptive, query-specific Retrieval And Shot selection), a method that improves TFM efficiency by selecting optimal shots based on local neighborhood analysis within the training set. Our results demonstrate that ARASH reduces the token length and memory usage of TFMs by $4.6\times$ and $4.9\times$, respectively, while preserving accuracy.

Martin Pham

Masters student (with Maryam Dehnavi)

Computer Science, University of Toronto

Martindopham@gmail.com

WebAssembly for browser-based scientific computing

We present two applications of WebAssembly for browser-based scientific computing. We first present a comparison of the Brian library extension Brian2WASM to a Rust-based spiking neural network implementation compiled to WebAssembly. We discuss the limitations to this approach and demonstrate some possible solutions using Rust closures. Secondly, we present a WebAssembly module for hyperdimensional computing on the Xiangqi (Chinese chess) board game. Hyperdimensional computing, sometimes called vector symbolic architectures, is a framework for representing data as algebraic operations in a high dimensional vector spaces (on the order of tens of thousands). We demonstrate preliminary results for board state encoding from Forsyth–Edwards notation and discuss next steps for implementing a vector symbolic game engine/artificial player. Together, these two examples demonstrate how WebAssembly can be deployed on the browser for scientific computing.

Alireza Razavi

Ph.D. student (with Masayuki Yano)
Aerospace, University of Toronto
a.razavi@mail.utoronto.ca

Registration-based Nonlinear Model-Order Reduction for Transonic Aerodynamics

Model-order reduction (MOR) is a class of methods for the rapid and reliable solution of parametrized partial differential equations (PDEs) in many-query settings, such as optimization and uncertainty quantification. MOR achieves computational reduction through offline-online decomposition: in the offline stage, we invoke the expensive full-order model (FOM) for select parameter values to learn the low-dimensional structure of the problem; in the online stage, we invoke the cheap reduced-order model (ROM) to efficiently solve the PDEs for many parameter values. Linear MOR enables accurate and efficient solution of many PDEs, on the order of 1% error and 1000x speedup relative to the FOM. However, for certain PDEs such as those in transonic aerodynamics, the parametric solution manifolds have spatially varying discontinuities, and linear MOR methods are unable to accurately approximate these solution manifolds in low-dimensional spaces. This is known as the Kolmogorov barrier. To overcome this barrier, we present a registration-based nonlinear MOR method which is scalable to large-scale problems. Our method introduces a parametrized, geometrically warped PDE, which enables us to compute a warped solution manifold where solution discontinuities are aligned. This manifold is then amenable to standard linear MOR techniques. We present the application of this method to large-scale, 2D and 3D transonic problems governed by the Euler and Reynolds-averaged Navier-Stokes equations.

Khan Enaet Hossain

Ph.D. student (with Dong Liang)
Mathematics and Statistics, York University
enaet09@yorku.ca

The Parallel DE-PDE Multicomponent Pollution Optimal Control Approach for Contamination Flows in Porous Media

In this study, we develop a problem-specific parallel DE-PDE optimization algorithm for efficiently solving the computationally intensive optimization problem arising from multicomponent groundwater contamination control. The proposed parallel DE-PDE algorithm tightly integrates the DE and PDE algorithms in parallel and exploits its population-based structure by distributing candidate control variables, such as injection rates, across multiple processors over the Message Passing Interface (MPI). We employ a stable and accurate PDE numerical solver to simulate multicomponent groundwater flow and contaminant transport in the DE-PDE approach. We consider the objective function to consist of the concentration-matching term and the piecewise-defined abatement cost function. A couple of numerical experiments validate the effectiveness of the parallel DE-PDE optimization algorithm for the multicomponent groundwater contamination control model.

Kimon Fountoulakis

Computer Science, University of Waterloo
kfountou@uwaterloo.ca

Complexity of Classical Acceleration for L1-Regularized PageRank

In this presentation, we discuss an open problem: how to bound the total work of practical accelerated methods for L1-regularized PageRank and when they can truly outperform non-accelerated methods. We focus on classical one-gradient-per-iteration accelerated methods, which are simple and practical but especially hard to analyze in this setting. We also emphasize that GPT-5.2 Pro, under close human guidance, generated all the proofs and much of the analysis, helping to identify a new trade-off and conditions under which acceleration helps, while the human researchers set the goals, enforced the right constraints, caught missing assumptions, and selected the final results.

Gurpreet Jagdev

Ph.D. student (with Na Yu and You Liang)
Mathematics, Toronto Metropolitan University
gjagdev@torontomu.ca

Coherence resonance in motif-embedded neural networks

Motifs, or small recurring connectivity patterns, are important mesoscale building blocks of complex biological networks. A central question is how such local structure shapes global dynamics in large noisy neuronal systems. Through the framework of coherence resonance, this work examines how motifs influence synchronization, coherent spiking, and network responses to noise. It further considers how these effects depend on the interplay between local connectivity patterns and overall network organization. By comparing motif-structured networks with suitable null models, the study shows that motifs can reorganize collective dynamics and produce more synchronous networks. Broadly, this work helps clarify how mesoscale organization contributes to emergent behavior in stochastic dynamical systems. Collaborators/Supervisors: Na Yu

Guilherme Macieira de Araujo

Ph.D. student (with Hans De Sterck)

Applied Mathematics, University of Waterloo

ghmaciei@uwaterloo.ca

Parallel-in-iteration optimization using multigrid reduction-in-time

Standard gradient-based iteration algorithms for optimization, such as gradient descent and its various proximal-based extensions to nonsmooth problems, are known to converge slowly for ill-conditioned problems, sometimes requiring many tens of thousands of iterations in practice. Since these iterations are computed sequentially, they may present a computational bottleneck in large-scale parallel simulations. In this work, we present a "parallel-in-iteration" framework that allows one to parallelize across these iterations using multiple processors with the objective of reducing the wall-clock time needed to solve the underlying optimization problem. Our methodology is based on re-purposing parallel time integration algorithms for time-dependent differential equations, specifically multigrid reduction-in-time (MGRIT), motivated by the fact that optimization algorithms often have interpretations as discretizations of time-dependent differential equations. We numerically demonstrate the efficacy of our approach on two model problems, a standard convex quadratic problem and the nonsmooth elastic obstacle problem, for which we observe fast MGRIT convergence analogous to its prototypical performance on partial differential equations of diffusion type. Some theory is presented to connect the convergence of MGRIT to the convergence of the underlying optimization algorithm. Theoretically predicted parallel speedup results are also provided. Coauthors: H. De Sterck, O. Krzysik.

Zelalem Arega Worku

Postdoctoral fellow (with David C. Del Rey Fernandez)

Applied Mathematics, University of Waterloo

zelalem.worku@uwaterloo.ca

On the Potential and Limits of High-Order Residual Correction Discretizations

Entropy-stable high-order discretizations are attractive for nonlinear conservation laws, but standard Hadamard-form two-point flux based formulations can be significantly more expensive than divergence-form schemes. As a cheaper alternative, the residual correction method adds a design-order correction that enforces conservation and energy or entropy stability. In this talk, I consider residual-correction discretizations based on continuous summation-by-parts (C-SBP) operators and focus on two questions that are central to their practical use: long-time accuracy and robustness. For model advection problems, I show that these discretizations satisfy an a priori error bound with at most linear growth in time when the SBP property holds, while departures from the SBP property can lead to faster error growth. I also discuss how the C-SBP framework helps improve the robustness of residual-correction discretizations while retaining much of their efficiency advantage. Numerical examples will be shown to illustrate these results.