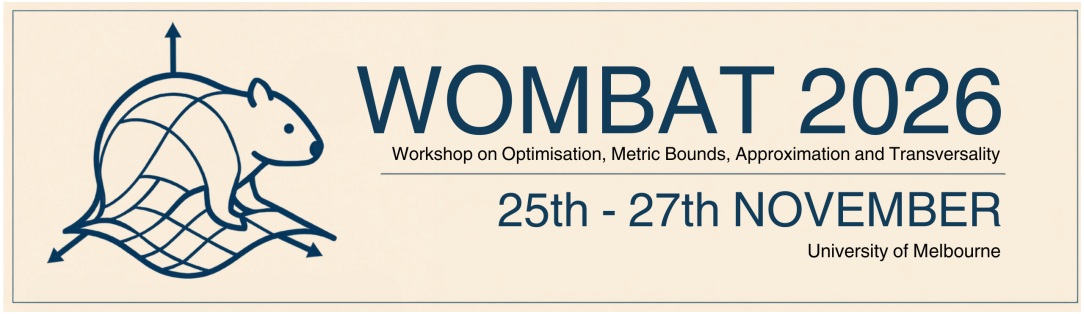


Workshop on Optimisation, Metric Bounds, Approximation and Transversality (WOMBAT)

joint with

SigmaOpt Workshop

The University of Melbourne

25–27 November 2026

Supported By

Australian Research Council

Mathematics of Computation and Optimisation, AustMS Special Interest Group

School of Mathematics and Statistics, The University of Melbourne

SigmaOpt, ANZIAM Special Interest Group



Conference Information

We acknowledge the Traditional Owners of the unceded land on which we meet: the Wurundjeri Woi-wurrung people. Aboriginal and Torres Strait Islander people are this country's first scientists, with deep and enduring knowledge of the land, waters and skies. We pay our respects to Elders, past and present, and we welcome all First Nations people here today. In making this acknowledgement, we commit to respectful and responsible conduct towards all others according to the Traditional lores of this land, particularly at times of formal ceremony.

In 2026, the annual Workshop on Optimisation, Metric Bounds, Approximation and Transversality (WOMBAT 2026) will be run in conjunction with the annual SigmaOpt workshop. Together, these workshops aim to bring together Australian and international researchers interested in all areas of optimisation.

Location

The conference will be held at The University of Melbourne's Parkville campus, in the Evan Williams Theatrette, room G03 (ground floor) of the Peter Hall Building (Building 160).

- University map
- Google Maps

Room Information

All talks will be held in Evan Williams Theatrette, room G03 (ground floor). See the building map later in the program for detailed location information.

The room has HDMI and USB-C ports to connect to projectors, or you can use a USB drive to load a presentation onto the room's desktop computer.

Public Transport

The University of Melbourne's Parkville campus is located just north of the Melbourne CBD, approximately 10min by tram or a 20min walk from Melbourne Central.

Tram The University of Melbourne/Swanston St (#1) tram stop is directly outside the Peter Hall Building. Trams to/from the Melbourne CBD are very frequent, including routes 1, 3, 5, 6, 16, 64, 67 and 72 that all go through Swanston St in the CBD. Alternatively, the route 19 tram goes along Elizabeth St to Royal Melbourne Hospital–Parkville Station/Royal Parade (#10) tram stop, which is a 10 min walk from the Peter Hall Building.

Train The closest train station is Parkville, on the Cranbourne, Pakenham and Sunbury line, and is on the edge of the campus. Alternatively, you can take a train to Melbourne Central, and then change to one of the Swanston St trams to campus.

The Transport Victoria journey planner website can give more specific transport advice. Public transport in Melbourne uses the myki card system, but you can also tap on and off with a contactless bank card or digital wallet. More information is available on the Transport Victoria website. Please note that contactless payments are currently for full-fare tickets only.

Airport Transfers

From Melbourne Tullamarine Airport (MEL), you can take the Skybus Melbourne City Express to the coach station at Southern Cross station in the CBD (approx. 30 mins). To get to campus, you catch the 86 or 96 tram from Spencer St/Bourke St (#1) to Swanston St/Bourke St (#6), then change to one of the Swanston St trams. Taxis and rideshare services like Uber are also readily available at the airport.

From Melbourne Avalon Airport (AVV), the https://www.skybus.com.au/avalon-city-express/ goes to the coach station at Southern Cross station, and takes approx. 90 mins. Taxis and rideshare services are also available.

Accommodation

There are many hotel options in the Melbourne CBD and next to campus in Carlton. The closest hotels to the conference venue are Crowne Plaza Melbourne Central (5 min walk) and the Travel Inn Hotel Carlton (10 min walk).

Tourism information

The nearest shopping and dining strip to the University of Melbourne is Lygon St in Carlton, a 5 min walk from the Peter Hall Building. The Melbourne CBD (20 min walk or 10 min tram) also has many shopping and dining options.

- Melbourne & Victoria tourism
- Melbourne weather forecast
- Transport Victoria (public transport information)
- Melbourne Airport transport

Catering

Morning and afternoon tea are provided in room G62 of the Peter Hall Building. There are many nearby options for lunch on-campus in buildings 162 and 168, adjacent to the Peter Hall Building, as well as numerous other options on Lygon St in Carlton (approx. 5 min walk).

Acknowledgements

Support for WOMBAT was provided by the Australian Research Council (DE240100006) and the University of Melbourne's School of Mathematics and Statistics.

WOMBAT is overseen by the Mathematics of Computation and Optimisation (MoCaO) Special Interest Group of the Australian Mathematical Society (AustMS), and the SigmaOpt workshop is overseen by the SigmaOpt ANZIAM Special Interest Group. Both MoCaO and SigmaOpt provided financial support for student travel.

Organising Committee

Lindon Roberts (The University of Melbourne)
Matthew Tam (The University of Melbourne)

Timetable

All talks will be held in Evan Williams Theatrette, room G03 (ground floor) of the Peter Hall Building (Building 160), and all times listed are AEDT (UTC+11). Tea breaks will be in the Maths Space room G06. Contributed talks should be 15 minutes long, plus 5 minutes for questions/changeover.

Day 1: Wednesday 25 November

Time	Speaker	Title
8:40 AM	*Organising committee*	*Registration/welcome*
9:00 AM	**Rachael Tappenden**	**WOMBAT Summer School**
10:00 AM	Jeffrey Christiansen	Adaptive Multi-Fidelity Benders Decomposition for AC Transmission Expansion Planning
10:20 AM	Srimali Gunasekara	Two-Stage Stochastic Programming under Endogenous Uncertainty for Shutdown Maintenance Scheduling
10:40 AM	Shahtaj Bano	Optimization of multiobjective chemotherapy treatment regimens
11:00 AM	*Tea break*	
11:40 AM	Thang D. Truong	Weak and strong convergence of a general splitting approach for monotone multioperator inclusions
12:00 PM	Hoa Bui	Exact Frank-Wolfe Method for the Inventory Routing Problem with a Logistic-Ratio Objective
12:20 PM	Jordan Collard	Extrapolation for Douglas–Rachford on Affine Subspaces
12:40 PM	Bethany Caldwell	Optimal control duality and the Douglas–Rachford algorithm
1:00 PM	*Lunch*	
2:00 PM	Andrew Eberhard	A Frank Wolfe Progressive Hedging Approach to Multi-objective Convex mixed integer stochastic optimisation
2:20 PM	Peter Gill	The Minimum Duration Path Problem: Exploiting Slopes
2:40 PM	Edward Lam	Conflict Analysis in Branch-and-Price
3:00 PM	Nikolaos Chatzis	Dispatch scheduling for hybrid renewable-hydrogen plants with non-linear constraints
3:20 PM	*Tea break*	
4:00 PM	Julien Ugon	A descent algorithm for best continuous polynomial spline approximation with free knots.
4:20 PM	Nadia Sukhorukova	Inverse problems through rational approximation
4:40 PM	Janosch Rieger	Galerkin approximations to the space of convex bodies by polytopes in nondegenerate V-representation
5:00 PM	*Day 1 ends*	

Day 2: Thursday 26 November

Time	Speaker	Title
9:00 AM	**Rachael Tappenden**	**WOMBAT Summer School**
10:00 AM	Chandula Fernando	A Multiobjective Model-Based Derivative-Free Trust-Region Method with Convergence and Complexity Guarantees
10:20 AM	Yuen-Man (Mandy) Pun	Regularized Projection Method for Randomized Zeroth-Order Optimization
10:40 AM	Amir Ali Farzin	Spherical-Design Oracles for Zeroth-Order Optimisation: Deterministic, Variance-Free Gradient Estimation
11:00 AM	*Tea break*	
11:40 AM	Ye Pu	Escaping Saddle Points in Gradient-Based Distributed Optimization via Random Initialization and Noise
12:00 PM	Jing Fu	Weakly-Coupled Restless-Bandit Models: Exponential Convergence in Probability
12:20 PM	Ziyue Zhao	Near optimal error bound of anchored spectral method for rigid motion synchronization
12:40 PM	Nam Ho-Nguyen	Margin-Based Thresholds for Learnability in Online Strategic Classification
1:00 PM	*Lunch*	
2:00 PM	**Haihao Lu**	**Plenary: GPU-Accelerated First-Order Methods for Linear Programming and Beyond**
3:00 PM	Andrew Calcan	Smooth approximation techniques and primal-dual methods
3:20 PM	*Tea break*	
4:00 PM	Lien Thuy Nguyen	Adaptive Stochastic ADMM for Nonconvex Optimisation
4:20 PM	Haoming Qu	A Proximal Point Algorithm for Compositions of Monotone Operators
4:40 PM	Liam Timms	A fast proximal-gradient algorithm for decentralised optimisation
5:00 PM	*Day 2 ends*	

Day 3: Friday 27 November

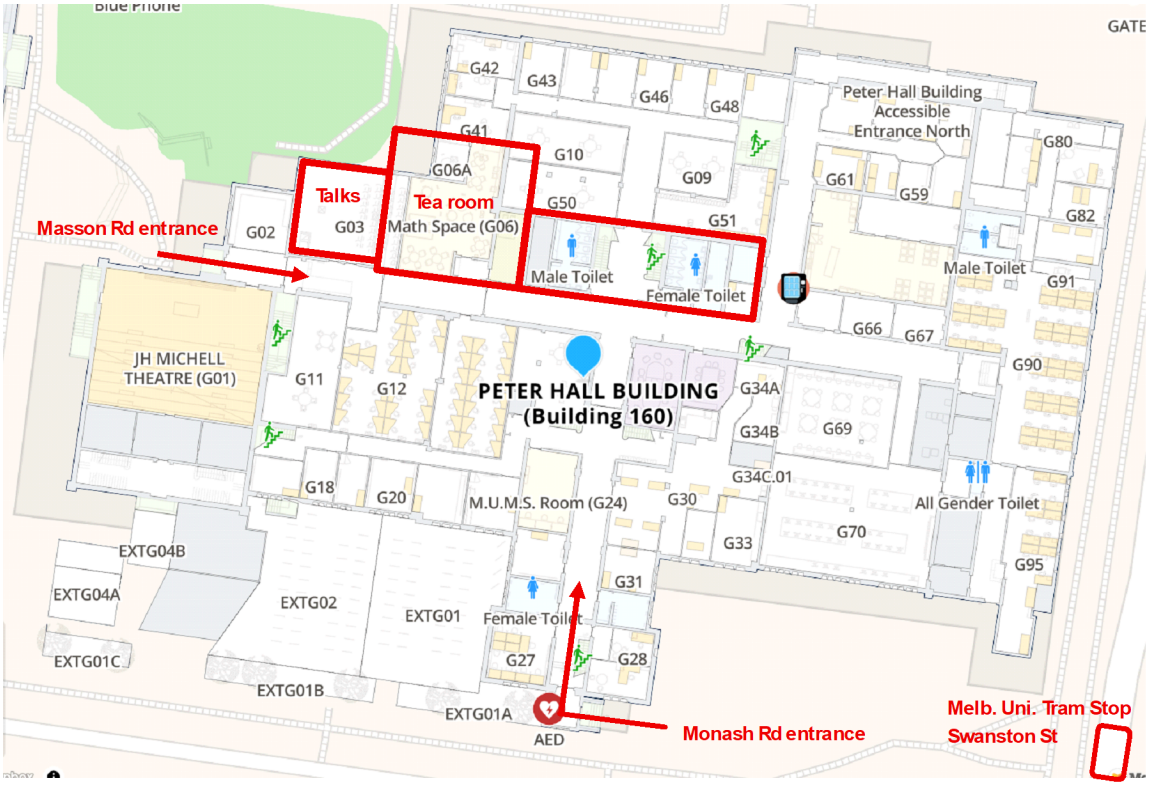
Time	Speaker	Title
9:00 AM	**Rachael Tappenden**	**WOMBAT Summer School**
10:00 AM	Gabriel Jarry-Bolduc	On the value of high-order accurate simplex gradients
10:20 AM	Gabriel Provencher Langlois	Exact algorithms for solving lasso and total variation problems via differential inclusions
10:40 AM	Zhengchao Wang	Overcoming Demand Censoring through Randomization for Risk-Sensitive Newsvendor Problems
11:00 AM	*Tea break*	
11:40 AM	Chefi Triki	Green Supplier Selection and Order Assignment: A Combinatorial Auction Framework
12:00 PM	Muhammad Abdullahi	TBC
12:20 PM	Alberto Marcolongo	The competitive two-echelon location routing problem
12:40 PM	Li Chen	Robust Bilevel Optimization Under Behavior Uncertainty
1:00 PM	*Lunch*	
2:00 PM	Matthew Hough	Convergence of the Frank-Wolfe Algorithm for Monotone Variational Inequalities
2:20 PM	Santanu Soe	GRAAL for variational inequalities revisited: sharp parameter range and adaptive step sizes
2:40 PM	Minh N. Dao	Primal-dual splitting for structured composite monotone inclusions with or without cocoercivity
3:00 PM	*Conference ends*	

Building Map

WOMBAT 2026 will be in the Peter Hall Building (building 160) of the University of Melbourne's Parkville campus.

- All talks will be in the Evan Williams Theatrette (ground floor, room G03)
- Tea breaks will be in the Maths Space room (ground floor, room G06)

If you enter the Peter Hall Building from the main entrance on Masson Rd, room G03 is the first room on your left (G02 on the map below is a waiting area). If you enter from the side entrance on Monash Rd (e.g. coming from the Melbourne University tram stop on Swanston St), follow the corridor straight ahead and then to your left; the room will be on your right.



Invited Speakers

WOMBAT Summer School

Leveraging curvature information in first order optimization algorithms

Rachael Tappenden
University of Canterbury

First-order optimization algorithms are a practical choice for many problems due to their simplicity, low computational and memory requirements, and ability to scale to high-dimensional settings. Nevertheless, their performance can deteriorate on ill-conditioned problems, they can struggle to navigate nonconvex landscapes, and they may exhibit slow escape from saddle points.

These talks will explore some of the challenges faced by first-order methods and investigate whether leveraging curvature information can improve their practical performance. In particular, we will consider strategies for obtaining partial or approximate curvature information at a low computational cost, and examine how this information can be exploited to improve the behaviour of first-order optimization algorithms.

Plenary Talk

GPU-Accelerated First-Order Methods for Linear Programming and Beyond

Haihao Lu
Massachusetts Institute of Technology

Can first-order methods and GPUs solve traditional mathematical optimization problems at scale? For decades, this seemed unlikely: conventional wisdom held that reliable LP solvers depended on simplex or interior-point methods and sparse factorizations, while first-order methods were considered too slow or inaccurate. This talk traces the PDLP research line from CPU-based PDLP through GPU-based cuPDLP to cuPDLPx. I will explain how GPU-friendly sparse matrix-vector operations, restarted primal-dual algorithms, and careful numerical design together achieve high accuracy, robustness, and competitive end-to-end performance. I will also discuss theoretical guarantees for linear convergence and infeasibility detection. Computational results cover standard LP benchmarks and practical applications, including problems with hundreds of millions of variables or nonzeros. Beyond its academic contributions, this line of research has helped drive a broader shift toward GPU-based first-order methods across the optimization solver ecosystem. These methods complement simplex and interior-point algorithms and become especially compelling at very large scales. Time permitting, I will conclude with extensions to convex quadratic and semidefinite programming.

Contributed Talks

Adaptive Multi-Fidelity Benders Decomposition for AC Transmission Expansion Planning

Jeffrey Christiansen
The University of Melbourne

Given an AC power network, the optimisation problem of finding the lowest-cost power dispatch plan which satisfies demand is non-convex and NP-hard. Transmission Expansion Planning introduces discrete variables corresponding to individual investment decisions, resulting in an even more difficult problem. Previous work on this problem has typically focused on heuristics, or on simpler approximations of the AC power flow equations. This talk will describe the Adaptive Multi-Fidelity Benders Decomposition algorithm, which converges to an optimal solution of the AC Transmission Expansion Planning problem. The algorithm also efficiently handles security constraints which ensure resilience in case of equipment failures.

Two-Stage Stochastic Programming under Endogenous Uncertainty for Shutdown Maintenance Scheduling

Srimali Gunasekara
Curtin University

This presentation outlines the research on optimising preventive maintenance planning during planned shutdowns in asset-intensive industries. These shutdowns are tightly constrained by fixed time, resource, and budget limits, requiring the selection of maintenance tasks. As not all tasks can be completed, some must be deferred, increasing the risk of future equipment failure and costly corrective maintenance. This problem is formulated as a two-stage stochastic program and shows that exact methods become intractable for realistic instances due to exponential growth in decision combinations and failure scenarios. To address this, I propose heuristic and approximation approaches, including a greedy first-stage framework with dominance-based screening, an adapted Sample Average Approximation (SAA) method, and scenario clustering with proxy recourse evaluation. The framework is evaluated on a benchmark dataset, with computational results demonstrating that it produces high-quality solutions within practical time limits. The presentation will focus on model formulation, algorithm design, and computational results.

Optimization of multiobjective chemotherapy treatment regimens

Shahtaj Bano
RMIT University

Chemotherapy treatment planning is not simply about finding the right drug dose. A treatment that improves tumor reduction may also increase side effects, cost, or the risk of relapse. In this work, we

develop a multiobjective optimization framework that brings these competing factors together and explores the trade-offs between them. We compare two optimization approaches, a scalarization-based method and NSGA-II, to identify different treatment regimens and examine how the choice of treatment changes when multiple objectives are considered simultaneously. This provides a foundation for moving from a single treatment goal toward more flexible and comprehensive chemotherapy planning.

Weak and strong convergence of a general splitting approach for monotone multioperator inclusions

Thang D. Truong
RMIT University

In this paper, we study a general forward-backward splitting approach for monotone inclusion problems that simultaneously accommodate set-valued, monotone Lipschitz continuous, and cocoercive operators. Weak convergence is established under a general setting. Strong convergence is obtained when both the maximally monotone operators are linear relations and the monotone Lipschitz continuous and cocoercive operators are linear. When the set-valued operators are normal cones to closed linear subspaces, we derive an explicit characterization of limit points of generated sequences using a new technique. Numerical experiments on a quadratic minimization problem with multiple-set split feasibility constraints are presented to illustrate the theoretical results.

Exact Frank-Wolfe Method for the Inventory Routing Problem with a Logistic-Ratio Objective

Hoa Bui
Curtin University

We propose an exact Frank-Wolfe method for the inventory routing problem with logistic ratio. The linear fractional objective is quasilinear, and convexification of the integer feasible set preserves the optimal value. We show that Frank-Wolfe generates a strictly improving sequence of integer solutions and reaches a global optimum in finitely many iterations. Computational experiments on standard benchmark instances show that Frank-Wolfe outperforms the exact algorithm of Archetti et al. (2019), certifying optimality on more instances while requiring substantially less computation time on instances solved by both methods. These results establish Frank-Wolfe as a simple and practical exact method for the routing problem with logistic ratio. In the end of the talk, we will go beyond the linear fractional setting and exploring conditions under which Frank-Wolfe identifies an optimal face in finitely many iterations for general nonlinear optimisation problems.

Extrapolation for Douglas–Rachford on Affine Subspaces

Jordan Collard
Curtin University

The Douglas–Rachford method can converge slowly, motivating extrapolation from successive iterates. For two intersecting affine subspaces, we show that the differences between successive iterates span a space whose dimension depends on the principal-angle geometry and the initial point, and can be much smaller than the ambient dimension. We prove that the Lyapunov surrogate method recovers the

Douglas–Rachford limit exactly using an extrapolation order equal to the dimension of this space. We also show that the Lyapunov surrogate method and GMRES applied to the Douglas–Rachford fixed-point equation use the same affine search spaces. We then examine how numerical conditioning affects recovery in practice.

Optimal control duality and the Douglas–Rachford algorithm

Bethany Caldwell
The University of Queensland

The Douglas–Rachford (DR) algorithm is a popular and well-studied algorithm but despite its prevalence there have been few studies applying this to infinite dimensional problems. We apply this algorithm to an infinite dimensional optimization problem, specifically, the dual of a linear-quadratic control problem with control constraints. We derive the fixed point of the DR operator as applied to this problem and using this expression for the fixed point we construct an optimality check to verify our numerical solutions. We propose an algorithm for our optimal control problem that gives both the primal and dual sequences and apply this to two example optimal control problems.

A Frank Wolfe Progressive Hedging Approach to Multi-objective Convex mixed integer stochastic optimisation

Andrew Eberhard
RMIT University

In this talk we will look at an extension of the Frank-Wolfe-Progressive Hedging (FWPH) Algorithm to the context of convex mixed integer stochastic optimisation (SMIP). We begin by looking at two stage SMIP and how we can develop a lower bound that converges to the Lagrangian relaxion dual bound to the problem. Convergence analysis is developed and we discuss how the algorithm may generated excellent primal solutions. We will then discuss the extension to multi-stage problems. A application to chemotherapy drug dosage optimisation will be discussed.

This is joint work with Shahtaj Bano and Minh Dao.

The Minimum Duration Path Problem: Exploiting Slopes

Peter Gill
Monash University

The minimum duration path problem (MDPP) is a time-dependent variant on the classic single-pair shortest path problem on a directed network. Each arc has a continuous piecewise-linear travel time function, rather than a static travel time, and the objective is to minimise the total travel duration over all "timed-paths" from the source to the target. Some applications of the MDPP arise in public transit routing, vehicle routing, ridesharing, and navigation. Previous work by He et al. (2022) [1] presented a "dynamic discretisation discovery" (DDD) method for the MDPP that may be viewed as a branch-and-bound algorithm. We build upon this method by recasting it into the framework of a mixed-integer program. We then obtain stronger intermediate lower bounds throughout the method based on classic linear programming theory and Lagrangian relaxation, allowing us to exploit "non-zero slopes" in the problem data (unlike, we claim, the classic DDD method).

References: [1] He, E. Y., Boland, N., Nemhauser, G., & Savelsbergh, M. (2022). Dynamic discretization discovery algorithms for time-dependent shortest path problems. INFORMS Journal on Computing, 34(2), 1086-1114.

Conflict Analysis in Branch-and-Price

Edward Lam
Monash University

Branch-and-price, like many tree search algorithms, suffers from repeatedly making the same branching decisions that lead to infeasibility or suboptimality. In contrast, propositional satisfiability and constraint programming have largely resolved this issue through conflict-driven clause learning and conflict analysis, which learn constraints, called nogoods, that prevent the same failure from recurring. Branch-and-price currently lacks conflict analysis, even though the pricing problem already certifies infeasibility or optimality of the master problem when it fails to find a column with negative Farkas cost or reduced cost respectively. Instead of discarding this evidence, this paper exploits it in a conflict analysis framework that identifies a subset of the branching decisions responsible for the non-existence of improving columns in the pricing problem, and then refutes them using a combinatorial Benders cut in the master problem. Experimental results on a multi-commodity flow problem demonstrate that enabling conflict analysis achieves a 2.5 times geometric mean speed-up on jointly solved instances, explores 9.8 times fewer branch-and-bound nodes in total, and more than doubles the number of instances solved.

Dispatch scheduling for hybrid renewable-hydrogen plants with non-linear constraints

Nikolaos Chatzis
Curtin University

Integration of hydrogen production in renewable plants is a relevant topic in regard to decarbonization and the push for net zero emissions. The option to divert electricity away from the market and into hydrogen provides a relief from risks related to the unpredictable nature of renewable sources. Implementation and optimal operation, not only it provides economic incentives for investors, but it also provides a better allocation of resources for the whole society. This talk aims to formulate the dispatch problem from first principles and showcase the challenges that are brought by the non-linear relation between energy and hydrogen production. Several modelling approaches are considered on how to treat the non-linear constraint including different styles of approximation or Lagrangian relaxation of the constraint. Particular, emphasis is placed for the conditions that allow the problem to be decomposable across time periods. This reduces the original problem to a series of independent subproblems that can be solved quite efficiently. When such decomposition is not allowed, it will be an outcome of storage constraints coupled across time. In that scenario, we will showcase techniques to bound the objective function and hence make inferences from the former to the latter version of the problem.

A descent algorithm for best continuous polynomial spline approximation with free knots

Julien Ugon
Deakin University

Approximating functions with a continuous piecewise polynomial is a challenging problem that requires solving a nonconvex problem. In this talk we will consider the Chebyshev approximation

problem, whose goal is to minimise the maximal deviation between the function and its approximation on an interval. We reformulate this problem as a bilevel programming problem and develop a descent algorithm that converges towards a stationary solution.

Inverse problems through rational approximation

Nadia Sukhorukova
Swinburne University of Technology

Inverse problems have numerous application in industry. There are many publications on this issue, including several textbook. Despite the significant amount of work, there is a gap between the mathematical theory and its applicability to real-life problems. In this paper, we consider the inverse problems as an approximation problem and apply optimisation methods to solve it. We use rational approximation due to its accuracy and relatively inexpensive algorithms. On the other hand, rational approximations can be used for implicit approximation of the level and sublevel sets as polynomial varieties. we run our numerical experiments on two types of problems: the problems where the level and sublevel sets are known (model validation) and tests datasets.

Galerkin approximations to the space of convex bodies by polytopes in nondegenerate V-representation

Janosch Rieger
Monash University

We introduce finite-dimensional approximations to the space of convex bodies based on polytopes in vertex representation. For a prescribed set of directions, the admissible point configurations form a polyhedral convex subcone of the Euclidean vector space, described by a finite system of linear inequalities. The interior of this cone contains only nondegenerate representations, in which all parameter points are distinct vertices of the represented polytope. We study the geometry of the parameter cone, redundancies in its defining inequalities, and natural projections of convex bodies onto the resulting polytope spaces. We derive quantitative approximation estimates in terms of how densely the prescribed directions cover the unit sphere and construct nested Galerkin sequences whose approximation error converges locally uniformly to zero. Finally, we use these approximation properties to establish the convergence of finite-dimensional approximations of constrained global optimization problems in the space of convex bodies.

A Multiobjective Model-Based Derivative-Free Trust-Region Method with Convergence and Complexity Guarantees

Chandula Fernando
The University of Melbourne

Model-Based Derivative-Free Optimisation (MBDFO) is a branch of derivative-free optimisation in which a local model is substituted for the function where gradient information is unavailable. It is an area that has been developing to handle black-box objective functions in unconstrained cases, constrained cases, and cases with stochastic objective functions. However, existing work has largely been developed for single-objective problems, with limited extension to the multiobjective setting. Motivated by the problem of optimal large-scale statistical survey design, in this talk I will present a MBDFO trust-region method

for solving deterministic, unconstrained multiobjective problems, using linear and quadratic interpolation models, with supporting convergence theory and complexity guarantees.

Regularized Projection Method for Randomized Zeroth-Order Optimization

Yuen-Man (Mandy) Pun
The University of Melbourne

In this talk, we present a unifying zeroth-order (ZO) framework that recovers both the vanilla ZO method and the gradient sketching method SEGA as special cases. The method projects the finite difference observations onto the subspace spanned by randomly sampled directions via a regularized projection, while reusing past gradient estimate information scaled by a memory parameter. We find that the method usually performs best when past information is discarded and we provide a partial explanation for it. Then, we introduce a smoothed surrogate of the objective whose gradient is exactly the expectation of our gradient estimator. This surrogate inherits convexity and Lipschitzness of the objective. Based on these properties, we develop convergence results when the function is convex and strongly convex. Finally, numerical results show that the regularized projection method typically outperforms other randomized ZO methods across different settings.

Spherical-Design Oracles for Zeroth-Order Optimisation: Deterministic, Variance-Free Gradient Estimation

Amir Ali Farzin
The Australian National University

In this work, we develop a zeroth-order (ZO) method that estimates the gradient with minimal bias and variance. Instead of the random directions often used in the ZO literature, we consider central differences with directions given by a fixed spherical t-design, a finite point set that integrates every polynomial of degree at most t exactly. The oracle is deterministic and thus does not induce any variance. We also propose a general finite difference method in the higher-ordersmoothness regime and show the bias of both finite difference methods. We then show the convergence of the proposed ZO method via the input-to-state stability concept. Numerical experiments show the superiority of our proposed method over the existing random ZO methods.

Escaping Saddle Points in Gradient-Based Distributed Optimization via Random Initialization and Noise

Ye Pu
The University of Melbourne

Gradient-based methods are widely used in distributed optimization due to their simplicity. However, in non-convex settings, they face significant challenges in escaping saddle points. In this work we consider the distributed consensus optimization (DCO) problem and focus on two representative algorithms—Exact First-Order Algorithm (EXTRA) and Distributed Gradient Descent (DGD). By leveraging random initialization, we establish almost-sure second-order convergence guarantees for EXTRA, enabling passive escape from saddle points. Using random perturbations, we establish high-probability second-order convergence guarantees for DGD, enabling active escape from saddle points and higher efficiency. Numerical experiments confirm our theoretical findings.

Weakly-Coupled Restless-Bandit Models: Exponential Convergence in Probability

Jing Fu
RMIT University

We study a finite time horizon Markov decision process (MDP) consisting of several groups of multi-action finite-state restless bandit processes, which are identical within each group. The bandit processes into different groups can be rather different. The bandit processes are subject to multiple weakly coupled constraints on their state and action variables. In contrast to prior studies that considered only a few specific policies/algorithms, here, we study the behaviours of the general stochastic process and, most importantly, the design of policies that guarantee its convergence to an ideal trajectory as the problem size increases. We prove that, for any policy in a rather general class, the resulting stochastic process converges in probability to a deterministic process as the system size (measured by the number of bandits) tends to infinity, at an exponential rate. Unlike the previous proofs, our exponential convergence does not rely on any non-degeneracy assumptions. It follows that the chosen policy asymptotically approaches optimality (with exponentially diminishing probability of suboptimality) in the size dimension if and only if the deterministic process coincides with optimality. As an example, in the multi-agent patrolling case, we propose a policy and prove that, in general, it converges in probability to optimality exponentially fast in the system size.

Near optimal error bound of anchored spectral method for rigid motion synchronization

Ziyue Zhao
The University of Hong Kong

The rigid motion synchronization problem aims to recover rigid motion group elements G_i^* for $1 \leq i \leq n$ from noisy pairwise measurements of $G_i^{*-1} \circ G_j^*$. This problem has a wide range of applications in signal processing, machine learning, and robotics, and has attracted considerable research attention in recent years. In this paper, we study the estimation performance of a proposed spectral method. Our contributions are twofold. First, we provide nearly optimal estimation guarantees for a unified noise model that incorporates missing data and outlier noise. Second, we introduce a hyperparameter β to improve the practical performance of the algorithm. We further conduct numerical experiments on synthetic data and the multiple-point set registration problem, demonstrating that our algorithm outperforms state-of-the-art spectral-type methods.

Margin-Based Thresholds for Learnability in Online Strategic Classification

Nam Ho-Nguyen
The University of Sydney

We study online strategic classification, in which the learner's deployed classifier influences the data it subsequently observes because agents can strategically alter their reported features. We consider a linear classification model with norm-based manipulation costs and assume that the agents' underlying true features satisfy a large-margin separability condition. Our goal is to characterize when meaningful learning remains possible under strategic behavior.

We show that learnability is governed by margin-based threshold phenomena. In particular, there are regimes in which no online method can guarantee finitely many prediction errors, and stricter regimes are

required to simultaneously control both prediction errors and strategic manipulation. To complement these limits, we develop a general reduction framework that converts standard online large-margin algorithms into algorithms for strategic classification. Using reported features alone yields guarantees in favorable regimes, while a proxy-data construction extends finite-mistake guarantees to the full learnable range.

Instantiations with perceptron and online maximum-margin methods recover and extend existing results, and show that stronger margin guarantees in the underlying learning algorithm translate directly into stronger guarantees in the strategic setting. We also study additive noise in reported features and show that the associated learnability thresholds depend sharply on the feedback model and on what information about the perturbation is available to the learner. Counterintuitively, learnability thresholds do not vary continuously with the noise level: the introduction of arbitrarily small additive noise can create a strictly harder learning problem than in the noiseless case, producing a discontinuous jump in the threshold for learnability.

Smooth approximation techniques and primal-dual methods

Andrew Calcan
Curtin University

In this talk, we investigate how principled approximation techniques interact with fundamental proximal algorithms and the methods they are dual to. This investigation is motivated by important problems in machine learning and inverse problems, where controlled approximation errors are often acceptable in exchange for improved numerical tractability and computational performance. We particularly examine the interplay between Moreau envelope smoothing and Fenchel duality. These insights are used to obtain local linear convergence rate bounds for the alternating direction method of multipliers. Conversely, smoothing on the dual problem motivates a novel algorithmic approach based on Tseng's alternating minimisation algorithm. We numerically validate this proposed approach on convex and nonconvex instances of the LASSO problem. These numerical results show our approach to consistently return acceptable solutions in less computational time than the fast iterative soft-thresholding and sparse relaxed regularised regression (SR3) algorithms for a certain class of linear predictive modeling problems.

Adaptive Stochastic ADMM for Nonconvex Optimisation

Lien Thuy Nguyen
RMIT University

In this work, we address a broad class of large-scale linearly constrained finite-sum difference-of-convex (DC) optimization problems arising in machine learning and signal processing. We propose a fast stochastic variance-reduced gradient alternating direction method of multipliers (FSVRG-ADMM) that incorporates a momentum-based variance-reduced gradient estimator to improve gradient tracking, with a generalized relaxation scheme and an adaptive residual-based relaxation rule. The proposed algorithm extends existing stochastic ADMM schemes and, under appropriate choices of the momentum and relaxation parameters, recovers standard ASVRG-ADMM as a special case. We further develop an abstract stochastic Kurdyka-Lojasiewicz (KL) framework for descent-type algorithms. Applied to FSVRG-ADMM, the framework establishes almost-sure whole-sequence convergence and expected convergence rates under separate stochastic KL conditions, while allowing the concave component of the DC term to be nonsmooth.

A Proximal Point Algorithm for Compositions of Monotone Operators

Haoming Qu
The University of Sydney

This talk presents a proximal-point-based splitting algorithm for compositions of monotone operators. We introduce a generalized reflected-resolvent framework and investigate its convergence properties under suitable assumptions. Numerical examples illustrate the proposed algorithm and its applications to nonlinear problems.

A fast proximal-gradient algorithm for decentralised optimisation

Liam Timms
The University of Melbourne

Consider a connected multi-agent network where agents coordinate to minimise a sum of convex functions, where each function has a smooth and nonsmooth component and is known only to a single agent. We propose a decentralised proximal-gradient-type algorithm for solving this problem, which we call Fast PG-EXTRA. In each iteration, agents evaluate the gradient of the smooth component of their local objective functions and the proximal operator of the nonsmooth component, followed by communicating the results with neighbouring agents. Fast PG-EXTRA can be viewed as an accelerated variant of the PG-EXTRA algorithm, in the sense that it enjoys $o(1/k^2)$ convergence rates in terms of the KKT residual and successive iterates, compared to PG-EXTRA's ergodic $O(1/k)$ convergence. Finally, we demonstrate the potential advantages of our algorithm with a numerical example for group LASSO regression.

On the value of high-order accurate simplex gradients

Gabriel Jarry-Bolduc
American University of Sharjah

The generalized simplex gradient is a first-order accurate approximation of the gradient, built from finite differences along an arbitrary set of sample directions. For smooth functions, its error expands in integer powers of the step size. We exploit this structure through Richardson extrapolation to obtain gradient approximations of any order of accuracy, at the cost of one additional set of function evaluations per order. We give an error bound that mirrors the classical first-order bound, and we determine the optimal step size when function values are noisy. We also prove that a centered variant uses the fewest function evaluations possible for a given order. In direct-search methods that halve the step size, the poll points of previous unsuccessful iterations may provide the extrapolated gradient at no extra cost, and we identify which order to use at each step size. Finally, we examine how the geometry of the sample directions affects accuracy. Numerical experiments identify some situations where high-order accuracy pays off.

Exact algorithms for solving lasso and total variation problems via differential inclusions

Gabriel Provencher Langlois
University of Illinois Urbana-Champaign

In this talk, I will present fast algorithms for solving the lasso problems and certain total variation minimization problems exactly, numerically up to machine precision, using novel differential inclusions techniques. Our proposed differential inclusions algorithms compute the solutions in finitely many steps and require no tuning of algorithmic parameters. I will present several numerical examples, including for post-processing numerical solutions of higher-order numerical methods for conservation laws.

Overcoming Demand Censoring through Randomization for Risk-Sensitive Newsvendor Problems

Zhengchao Wang
The University of Sydney

We study online generalized newsvendor learning under censored demand and convex nonlinear shortage costs. Demand may follow an arbitrary sequence, and the decision maker observes only sales censored by the chosen inventory level. Under a nonlinear shortage penalty, a stockout conceals the shortage magnitude and hence the subgradient required by a first-order learning method. We address this information problem through randomized tail probing: the implemented order is perturbed above a virtual base-stock level so that an unbiased stochastic subgradient can be recovered from censored observations. We characterize the class of admissible probing distributions and derive a regret bound that explicitly separates the statistical cost of gradient estimation from the operational cost of exploratory inventory. We then optimize the probing distribution with respect to this trade-off. For piecewise-linear shortage penalties, the optimal distribution is obtained through a linear-time recursion over the curvature knots; for penalties with absolutely continuous curvature, it is characterized by a scalar ordinary differential equation. The resulting policies achieve $O(T^{2/3})$ expected regret, with problem dependence summarized by a curvature-complexity measure. When the shortage penalty is linear, probing is unnecessary and the standard $O(\sqrt{T})$ regret rate is recovered.

Green Supplier Selection and Order Assignment: A Combinatorial Auction Framework

Chefi Triki
University of Kent

Outsourcing is increasingly adopted by firms aiming to improve operational efficiency and strengthen competitive advantage, which has intensified the importance of supplier selection and order allocation (SSOA) decisions. Although combinatorial auctions are widely regarded as an effective procurement mechanism, their application to multi-sourcing SSOA problems has received limited attention in the literature.

In this study, we introduce a new Winner Determination Problem formulation that extends existing models by enabling suppliers to bid on several bundles, including overlapping combinations. To solve the problem, we develop four novel heuristic procedures derived from the Kernel Search framework, incorporating both economic and environmental sustainability considerations into procurement decisions.

Computational experiments show that the proposed approaches perform effectively across a range of scenarios. In particular, the weighted frequency-based rarity and value-density heuristic provides the most favorable balance between solution quality and computational efficiency. The managerial analysis further indicates that auction-driven SSOA systems can generate considerable cost reductions relative to conventional approaches by exploiting economies of scope, while also emphasizing the trade-offs among supplier diversification, procurement costs, and risk exposure. Overall, this work contributes both practically and methodologically to the procurement and operations management literature, and opens avenues for future studies in stochastic and dynamic combinatorial auction environments.

Acknowledgments: Research reported in this paper was supported by the Qatar Research Development and Innovation Council [ARG01-0430-230029].

TBC

Muhammad Abdullahi
RMIT University

TBC

The competitive two-echelon location routing problem

Alberto Marcolongo
Aarhus University

Multi-echelon distribution networks are widely used in transportation systems due to their flex- ibility and efficiency. Their application ranges from logistics and supply chain management to urban logistics, where social considerations play an increasingly important role. In these contexts, a well-studied problem is the two-echelon location routing problem (2E-LRP), which simultaneously accounts for facility location and vehicle routing decisions. However, most studies assume that the set of customers to be served is known in advance and does not depend on competitive interactions between firms.

To address this limitation, we introduce the Competitive Two-Echelon Location Routing Problem (2E-CLRP), which integrates competitive customer capture to facility location and routing decisions. To the best of our knowledge, this is the first study that considers competitiveness within a two-echelon location routing framework. In this setting, each customer has a utility parameter associated with each facility location, and will choose the largest-utility location among the open ones. The utility parameter can reflect the service level a company offers depending on the location of the facility. We propose a flow-based formulation that adapts existing formulations for the two-echelon vehicle routing problem. We solve this model using exact solution methods that we evaluate on a new benchmark set of instances specifically designed for the 2E-CLRP. The solution methods tighten the formulation with problem-specific cuts, and reduce the problem size by exploiting structural observations on valuable facility locations and on the influence areas of facilities.

Preliminary computational experiments show that the proposed cuts and the problem-size reduction techniques reduce computational times and enable our approach to outperform a straightforward MILP formulation.

Robust Bilevel Optimization Under Behavior Uncertainty

Li Chen
The University of Sydney

We study a bilevel linear optimization problem in which the follower's objective coefficients are latent and must be inferred from observed decisions. We propose a type-I Kantorovich behavioral optimal-transport model that uses the suboptimality loss of each observed decision as the transportation cost, thereby linking distributional ambiguity directly to observed follower behavior. The resulting robust optimization model admits an exact measure-dual reformulation with a finite optimal multiplier for every feasible ambiguity radius, including the minimum feasible radius. We develop two scalable

solution approaches based on cost-scenario and follower-decision column-and-constraint generation, and characterize the precise conditions under which their separation problems admit mixed-integer linear programming reformulations. We further establish finite-sample guarantees for the proposed robust optimization framework and demonstrate its computational performance and practical robustness across applications in logistics and infrastructure planning.

Convergence of the Frank-Wolfe Algorithm for Monotone Variational Inequalities

Matthew Hough
University of Waterloo

We discuss the convergence of the Frank-Wolfe algorithm for monotone variational inequalities over compact convex sets and discuss the applications of this result. Under a monotone C^1 operator and vanishing, non-summable step sizes, we establish asymptotic convergence of the algorithm. In the strongly monotone case, the iterates converge to the unique solution, thereby proving Hammond's conjecture on generalized fictitious play. We also discuss barriers to proving convergence rates in the general setting and explain how stronger assumptions on the constraint set allow such rates to be obtained.

GRAAL for variational inequalities revisited: sharp parameter range and adaptive step sizes

Santanu Soe
The University of Melbourne

This paper revisits the Golden Ratio Algorithm (GRAAL) of Malitsky (Math. Program. 184, 383–410, 2020) for monotone variational inequalities. The original analysis restricts the convex-combination parameter c to $1 < c \leq \varphi$, where $\varphi = (1+\sqrt{5})/2$. Through a new Lyapunov analysis, we prove the convergence of GRAAL for every $1 < c < 3$. Further, a two-dimensional linear counterexample shows that this parameter range is sharp: for every $c \geq 3$ and every positive constant step size, the iterates can be unbounded. Building on this analysis, we propose a simpler version of the adaptive GRAAL (aGRAAL) that requires neither a line search nor a prescribed upper bound on the step sizes, and demonstrate its convergence for locally Lipschitz cost operators over the same parameter range. For both methods, we establish an $O(1/N)$ ergodic gap bound and, under a residual-based error bound condition, prove the R-linear convergence of the adaptive iterates. Finally, numerical experiments are conducted to compare the proposed strategies with the existing ones.

Primal-dual splitting for structured composite monotone inclusions with or without cocoercivity

Minh N. Dao
RMIT University

We propose a primal-dual splitting framework for a broad class of structured composite monotone inclusions that involve finitely many set-valued operators, compositions with bounded linear operators, and single-valued operators possibly without cocoercivity. Unlike standard product-space reformulation approaches, the proposed framework directly exploits the problem structure through some coefficient matrices, leading to explicit algorithms that are easy to customize for different computational requirements. It accommodates heterogeneous cocoercive or Lipschitz constants, different resolvent parameters,

providing substantial flexibility in algorithm design and distributed implementation. We introduce the notion of quasicomonotonicity and develop a unified convergence analysis covering both cocoercive and noncocoercive settings, yielding broader admissible stepsize ranges under weaker assumptions than recent methods. The framework unifies and improves several contemporary splitting methods while providing a systematic blueprint to generate new graph-based algorithms. We demonstrate the practicality of the approach by a numerical experiment on cancer detection using the decentralized fused LASSO problem.
