




WOMBAT 2025

Workshop on Optimisation, Metric Bounds, Approximation and Transversality

26th - 28th NOVEMBER

University of Queensland

WOMBAT 2025 Handbook

We are delighted to welcome you to the [Workshop on Optimisation, Metric Bounds, Approximation and Transversality](#) (WOMBAT) in 2025, held at the University of Queensland (UQ), Brisbane. This conference allows us to explore new ideas and share knowledge on any topic related to mathematical optimisation: theoretical, computational or applied.

Workshop Information

- **Venue:** Room 348, Physiology Lecture Theaters Building (Building 63), St Lucia campus [[Google Map](#)]
- **Date:** 26th - 28th November 2025
- **Time:** Main Conference Activities - 9am-5pm.

Getting to The University of Queensland

The University of Queensland's St Lucia campus is located just 7 km from Brisbane's city centre, nestled along the Brisbane River. It is easily accessible by public transport, including **buses**, **ferries** (**CityCat**), and **cycling paths**.

- **Bus:** Several bus routes travel directly to UQ St Lucia, including the **M2 "Metro Bus"** Route with stops in South Bank and the CBD. A number of other bus routes stop at the UQ Lakes and Chancellor's Place stops, the main drop-off points on campus.
- **Ferry:** The **CityCat ferry** service stops at the UQ St Lucia terminal, offering a scenic and convenient route from various parts of Brisbane.
- **Train:** The closest trains stations to UQ are **Toowong** and **Boggo Road**, which are a short bus ride to campus.
- **Cycling:** UQ is well-connected by bike paths, and there are bike racks and end-of-trip facilities available on campus.

From Brisbane Airport, you can take the **Airtrain** to the city (e.g., South Bank, Central Station or Roma Street Station), then transfer to a bus or ferry to reach UQ. Taxis and rideshare services like Uber are also readily available at the airport.

Public transport in Brisbane uses the **go card** system, but you can also tap on and off with a contactless bank card or digital wallet. More information is available on the [Translink website](#).

Travel & Health Information

If you're traveling internationally, please check the latest health and visa requirements for entering Australia.

If you need a letter confirming your participation or attendance at WOMBAT 2025, please contact wombatconference2025@gmail.com. **We recommend applying early if you require a visa.**

Accommodation

There are a variety of short-term accommodation options available both **on-campus** and **near the University of Queensland's St Lucia campus**, suitable for workshop attendees.

On-Campus Accommodation

- **UQ Res Short-Stays: Kev Carmody House** - One-bedroom units with ensuite bathrooms, social spaces, and a rooftop pool; **48 Walcott Street** - Two-bedroom apartments with shared kitchen and balcony, plus communal BBQ and recreation areas.
- **Residential Colleges:** During non-semester periods, some UQ colleges offer short-stay rooms with optional meal plans. These include [King's College](#), [International House](#) and [Grace College](#).

Off-Campus Accommodation

- **Essence Suites Taringa** - Offers a 15% discount for UQ visitors (use code "UQ" when booking).
- **Jephson Hotel & Apartments (Toowong)** - Boutique hotel with kitchenettes, close to shops and public transport.
- **Hotel Diana (Woolloongabba)** - Affordable, spacious rooms with free parking and a fitness centre.
- **Courtyard by Marriott (South Bank)** - Stylish rooms near the river and South Bank Parklands.
- **St Lucia Garden Apartments** - Budget-friendly and walkable to campus.
- [For more off-campus accommodation options](#)

Acknowledgments

We would like to acknowledge the generous support of the University of Queensland's [School of Mathematics and Physics](#), the University of Sydney's [Discipline of Business Analytics](#), and [Biarri](#). Their contributions have been instrumental in making this workshop possible and in fostering collaboration across academia and industry.

Organising Committee

- [Fred Roosta](#) (University of Queensland)
- [Oscar Smee](#) (University of Queensland)
- [Alexander Lim](#) (University of Queensland)

Timetable

Day One - Wednesday 26/11

Summer School	
09:00 - 10:00	Stephen Wright - Bilevel Optimization (Part I)
10:00 - 10:30	Morning tea
10:30 - 11:30	Stephen Wright - Bilevel Optimization (Part II)
Talks	
11:30 - 12:00	Iman Shames - Line-Search Filter Interior Point Differential Dynamic Programming for Optimal Control with Nonlinear Equality Constraints
12:00 - 13:00	Lunch break
13:00 - 15:00	Tomas Lagos - Data-Driven Optimization for Meal Delivery: A Reinforcement Learning Approach for Order-Courier Assignment and Routing at Meituan Vinesha Peiris - KKT-based optimality conditions for neural network approximations Lan-Phuong Nguyen - Efficient Online Large-Margin Classification via Dual Certificates Paul Grigas - Beyond Discretization: Learning the Optimal Solution Path
15:00 - 15:30	Afternoon tea
15:30 - 17:00	Lindon Roberts - Direct Search for Linearly Constrained Derivative-Free Optimization Fangyu Liu - Robust estimators for stochastic model-based DFO Thakshila Rajapaksha - From Linear to Superlinear: Advances in Derivative-Free Optimisation Algorithms
Conference Mixer	
17:00 - 18:30	Conference Mixer - Location TBD

Day Two - Thursday 27/11

Summer School	
09:00 - 10:00	Stephen Wright - Bilevel Optimization (Part II)
10:00 - 10:30	Morning tea
10:30 - 11:30	Stephen Wright - Bilevel Optimization (Part IV)
Talks	
11:30 - 12:00	Ponpot Jartnillaphand - Bilinear Logic-Based Benders Decomposition
12:00 - 13:00	Lunch break
13:00 - 15:00	Nam Ho-Nguyen - Convergence, Well-Posedness and Stochasticity in Convex Bilevel Optimization Tamanna yadav - Optimality analysis of bilevel programming problems and its application for the electric vehicle charging stations challenges Joseph Wilson - Cheap, Stable Probabilistic Sampling using First-Order Optimization Methods Oscar Smee - An inexact Primal-Dual Newton method for nonconvex optimization with equality constraints
15:00 - 15:30	Afternoon tea
15:30 - 17:00	Felipe Atenas - Variable stepsize splitting methods as relocated fixed-point iterations Thang D. Truong - A general approach to distributed operator splitting Thomas Chaffey - Circuit realization and hardware linearization of monotone operator equilibrium networks

Day Three - Friday 28/11

Plenary Talk	
09:00 - 10:00	Dmytro Matsypura - -
10:00 - 10:30	Morning tea
Talks	
10:30 - 12:00	Ayesha Sohail - Planning Renewable Energy Growth in Australia Using Adaptive Optimization Alba Olivares-Nadal - Optimizing Treatment Allocation to Maximize the Health of a Population Keith Moffat - Real-Time Power System Congestion Management using Feedback Optimization
12:00 - 13:00	Lunch break
13:00 - 15:00	Srimali Gunasekara - Shutdown Maintenance Scheduling using Two-Stage Stochastic Programming with Endogenous Uncertainty Johnson Phosavanh - Scheduling with cooperative agents Saeid Kazemzadeh - Multi-Objective Structural Optimization Using Evolutionary Algorithms: Challenges and Future Directions Kasra Khosoussi - Measure Once, Estimate Twice: Joint State and Noise Covariance Estimation
15:00 - 15:30	Afternoon tea
15:30 - 16:30	Li Chen - Robust Data-Driven CARA Optimization Nadezda Sukhorukova - General rational and neural networks approximation: combination and cross-pollination

Planary Speakers



Stephen Wright



Dmytro Matsypura

Summer School - Bilevel Optimization

Bilevel optimization is a nonlinear constrained optimization problem in which the constraints include a requirement that one of the variables is a solution of another, “lower-level” optimization problem. Such problems have many applications, particularly in machine learning, and algorithms to solve them have been the subject of much recent research. Even when the lower-level problem is strongly convex, the problem presents many challenges to algorithm design. (When the lower-level problem is *not* strongly convex, the objective may even be discontinuous.) In these talks, we start by surveying some applications of bilevel optimization and discussing the features of the problem that make it difficult to solve. We then discuss algorithms, including methods that make use of second derivatives, methods that use only first derivatives, and methods suitable for finite-sum objectives that require only unbiased estimates of first derivatives.

Talks & Abstracts¹

Optimizing Treatment Allocation to Maximize the Health of a Population

[Alba Olivares-Nadal](#), a.olivares_nadal@unsw.edu.au

Recent shifts in global health priorities have positioned Population Health Management (PHM) as a central area of focus. However, optimizing PHM strategies presents several challenges: managing high-dimensional patient covariates, tracking their evolution and long-term response to interventions, and accounting for the inflow and outflow of individuals within the population. In this paper, we propose a novel approach based on measurized Markov Decision Processes (MDPs) that integrates all of these components. Specifically, we consider a setting in which a treatment with population-level benefits is available but scarce, and model an MDP that optimizes the long-term distribution of the healthcare population subject to expected capacity constraints. This formulation allows us to bypass both the dimensionality and practical challenges of handling and tracking individual patient covariates across the population. To ensure ethical compliance, we introduce a non-maleficence constraint that limits the allowable mortality rate. To solve the resulting infinite-dimensional problem, we develop an approximate solution method that reduces the task to identifying a finite set of high-performing treated and untreated patients. Despite the complexity of the underlying structure, our approach yields a simple and clinically implementable index policy: a patient is selected for treatment if their adjusted impactability exceeds a specified threshold. The adjusted impactability captures the long-term consequences of receiving, or not receiving, treatment. While straightforward to apply, the policy remains flexible and can incorporate general machine learning models.

Planning Renewable Energy Growth in Australia Using Adaptive Optimization

[Ayesha Sohail](#), ayesha.sohail@sydney.edu.au

Australia is increasing its use of renewable energy but planning how quickly to build new solar and wind capacity remains a challenge. Building too fast can increase costs, while moving too slowly keeps the electricity system dependent on fossil fuels. This study presents a simple optimization approach to help choose renewable expansion plans that are both affordable and reliable. The model considers how electricity demand changes over time and how renewable output can vary with weather. It then uses an adaptive optimization method to test different build-rate decisions and identify strategies that keep total system costs low while ensuring the power supply is still stable. The results show that flexible, step-by-step planning performs better than fixed long-term schedules, especially when future demand is uncertain. This approach offers a practical tool to support energy transition decision-making in the Australian grid.

Robust estimators for stochastic model-based DFO

[Fangyu Liu](#), fliu0784@uni.sydney.edu.au

Derivative-free optimisation (DFO) methods are widely used when objective functions are expensive, noisy, or lack derivative information. In stochastic settings, model-based DFO algorithms typically rely on sample means to aggregate repeated noisy evaluations. However, under heavy-tailed or outlier-prone noise, mean-based estimators can become unreliable and require many samples. In this talk, I will present a robust alternative using median-based estimators within stochastic trust-region frameworks. I will show that the median achieves comparable accuracy to the mean in light-tailed noise, while providing significantly better robustness in heavy-tailed regimes, supported by theoretical error bounds and sample complexity analysis.

¹in alphabetical order.

Variable stepsize splitting methods as relocated fixed-point iterations

Felipe Atenas, felipe.atenas@unimelb.edu.au

Splitting methods exploit model structures to decompose complex problems into simpler pieces easier to handle. One prominent instance is the Douglas-Rachford algorithm, which stands out due to its simplicity and numerical stability. Traditional convergence guarantees assume constant stepsizes, while the theory with variable stepsizes is scarce. The fundamental challenge in varying this parameter stems from the stepsize-dependence of the fixed-point set of the Douglas-Rachford operator, preventing the use of classical arguments to guarantee convergence. To address this limitation, we propose a novel variant of the Douglas-Rachford algorithm that allows updating the stepsize between iterations, by composing the original Douglas-Rachford iteration with a ‘fixed-point relocater’ operator. For finding a zero of the sum of two maximally monotone operators in a Hilbert space, we show that the resulting relocated Douglas-Rachford method converges weakly to a solution of the monotone inclusion problem, under mild assumptions on the asymptotic behaviour of the sequence of stepsizes. We also examine the improvement in the numerical performance when using the relocated fixed-point variant, and extensions to distributed settings and to methods of forward-backward type.

Line-Search Filter Interior Point Differential Dynamic Programming for Optimal Control with Nonlinear Equality Constraints

Iman Shames, iman.shames@unimelb.edu.au

We present FilterDDP, a differential dynamic programming algorithm for solving discrete-time, optimal control problems (OCPs) with nonlinear equality constraints. Unlike prior methods based on merit functions or the augmented Lagrangian class of algorithms, FilterDDP uses a step filter in conjunction with a line search to handle equality constraints. We identify two important design choices for the step filter criteria which lead to robust numerical performance: 1) we use the Lagrangian instead of the cost as one of the filter criterion and, 2) for the stopping criteria and backward pass Hessians, we replace the value function gradient with an estimated dual variable of the constraints that characterise the dynamics. Both choices are rigorously justified, for 2) in particular by a formal proof of local quadratic convergence. In addition to providing a primal-dual interior point extension for handling OCPs with both equality and inequality constraints, we validate FilterDDP on three contact implicit trajectory optimisation problems which arise in robotics.

Scheduling with cooperative agents

Johnson Phosavanh, johnson.phosavanh@sydney.edu.au

We consider a new class of multi-agent scheduling problems where some jobs are allowed to be completed by more than one agent, each with a different processing time, allowing us to represent workers with different skill sets. This problem can be considered as a multi-stage optimization problem, where the assignment of jobs to agents needs to be decided, followed by the determination of an optimal schedule or vice versa. We show that this class of problems is not strictly easier or harder than the conventional competing or nondisjoint job set scenarios, and we provide an extensive analysis of problems consisting of two agents on a single machine with various scheduling criteria.

Cheap, Stable Probabilistic Sampling using First-Order Optimization Methods

Joseph Wilson, joseph.wilson1@student.uq.edu.au

Neural networks have shown impressive predictive results, but accurate quantification of the uncertainty in their predictions remains an open problem. Amongst promising research directions is the combination of Gaussian Processes (GP), the gold-standard machine learning method for epistemic model uncertainty quantification for small-scale problems, with the Neural Tangent Kernel

(NTK), a covariance function that determines the functional dynamics of a neural network during gradient descent training. However, forming uncertainty estimates from a GP naively incurs a $O(n^3)$ cost from inverting the Gram matrix, where n is the number of training points. Further, for the NTK, the spectrum of the Gram matrix contains many small eigenvalues, which affects the accuracy of the solution of the resultant linear system. In this talk, we show that by minimizing a strongly-convex optimization problem using (stochastic) gradient descent, we can efficiently compute samples from this NTK-GP, from whom the sample covariance gives a meaningful estimate of prediction uncertainty. Through this, we sidestep the scalability and stability issues of an NTK-GP.

Measure Once, Estimate Twice: Joint State and Noise Covariance Estimation

Kasra Khosoussi, k.khosoussi@uq.edu.au

In robotics and computer vision estimation problems it is typically assumed that the measurement noise covariance matrix is known. However, in practice, we often rely on ad hoc estimates (e.g., isotropic or diagonal) determined through (offline) trial-and-error tuning. This talk will cover a recent work (<https://arxiv.org/abs/2502.04584>) on joint estimation of the noise covariance matrix alongside primary parameters (such as poses and 3D points) online from measurements corrupted by Gaussian noise. We'll show that the joint problem exhibits a convex structure and provide a full (analytical) characterization of the optimal noise covariance estimate within joint maximum a posteriori and likelihood frameworks and several variants. Leveraging this theoretical result, we propose two novel algorithms that jointly estimate the primary parameters and the noise covariance matrix and provide convergence analysis. Finally, we'll show how this method can be used in the pose-graph optimization problem in robotics to estimate the noise covariance matrix online.

Real-Time Power System Congestion Management using Feedback Optimization

Keith Moffat, keith.moffat@unimelb.edu.au

As intermittent renewable generation and the demands on our electric infrastructure increase, grid operators are looking for new tools that can orchestrate fast-responding resources in real-time. One option is Feedback Optimization, which makes iterative adjustments to the power injections of controllable resources based on measurements from the physical system. Thus, the optimization algorithm is run in closed-loop with the physical system, rather than open-loop based on a model of the system, and naturally accounts for model mismatch. A feasibility study conducted with a German grid operator is determining how to use Feedback Optimization to provide grid operators with agile and intuitive control tools that reduce operation costs. The talk discusses the different optimization algorithms that can be brought to bear on the physical system, and concludes by discussing the pros and cons of each.

Efficient Online Large-Margin Classification via Dual Certificates

Lan-Phuong Nguyen, lanphuong.nguyen@sydney.edu.au

Online classification is a central problem in optimization, statistical learning and data science. Classical algorithms such as the perceptron offer efficient updates and finite mistake guarantees on linearly separable data, but they do not exploit the underlying geometric structure of the classification problem. We study the offline maximum margin problem through its dual formulation and use the resulting geometric insights to design a principled and efficient algorithm for the online setting. A key feature of our method is its translation invariance, inherited from the offline formulation, which plays a central role in its performance analysis. Our theoretical analysis yields improved mistake and margin bounds that depend only on translation-invariant quantities, offering stronger guarantees than existing algorithms under the same assumptions in favorable settings. In particular, we identify a parameter regime where our algorithm makes at most two mistakes per

sequence, whereas the perceptron can be forced to make arbitrarily many mistakes. Our numerical study on real data further demonstrates that our method matches the computational efficiency of existing online algorithms, while significantly outperforming them in accuracy.

Robust Data-Driven CARA Optimization

[Li Chen](#), li.chen@sydney.edu.au

Inspired by Von Neumann’s foundational work on expected utility, this paper extends the data-driven expected utility maximization framework under constant absolute risk aversion (CARA) preferences by incorporating nonlinear payoff functions and enhancing robustness. These improvements aim to broaden the applicability of utility maximization in prescriptive analytics and to improve out-of-sample performance, particularly in data-scarce settings, respectively. The combined nonlinearity of exponential utility and payoff functions presents significant computational challenges in ensuring robustness. We show that existing safe tractable approximation techniques are overly conservative, potentially causing unexpected infeasibility. To address this, we define a notion of consistent approximation and propose an augmentation technique that provably retains the consistency in the approximation. Theoretical insights are validated through computational experiments in two applications: data-driven portfolio optimization and facility location optimization. Our results demonstrate the effectiveness of incorporating robustness into prescriptive analytics, yielding solutions that significantly outperform traditional empirical and stochastic optimization. Our utility maximization framework is therefore not only robust, but also interpretable, scalable, and adaptable to a wide variety of real-world risk-aware decision-making problems involving nonlinear payoffs.

Direct Search for Linearly Constrained Derivative-Free Optimization

[Lindon Roberts](#), lindon.roberts@unimelb.edu.au

Direct search methods are a popular technique used for derivative-free optimization (DFO). A key benefit of direct search is its flexibility, being suitable for constrained, nonsmooth and mixed integer problems, for example. However, the existing theory for direct search with linear constraints does not allow steps to be taken towards nearby constraints, despite such steps being very helpful in practice. In this talk, I will outline a new convergence theory for direct search, how it justifies these practical choices, and explain the theoretical links to other DFO methods.

General rational and neural networks approximation: combination and cross-pollination

[Nadezda Sukhorukova](#), nsukhorukova@swin.edu.au

Neural networks (NNs) are popular techniques in modern science and engineering applications, where some type of approximation is required. These approaches are able to produce accurate approximations to nonsmooth and non-Lipschitz functions, including multivariate domain functions. Modern computer packages are designed in such way that the construction of approximations by NNs is straightforward for the users and they do not need to know the mathematics behind the scene. This seems to be convenient, but eventually most users would want to open the black-box and improve it. One possibility is to use rational approximation, which combines the high approximation accuracy and the simplicity of the optimisation tools: the algorithms are based on the repeated application of standard linear programming techniques that are part of most modern computer packages. Our numerical experiments demonstrate the efficiency of rational approximation, even when the number of approximation parameters (that is, the dimension of the corresponding optimisation problems) is small, while for neural networks a higher number of decision variables is required. We demonstrate our findings in numerical examples, one of them is approximating solutions of the KdV (Korteweg-de Vries) equation which appears in fluid dynamics.

Convergence, Well-Posedness and Stochasticity in Convex Bilevel Optimization

Nam Ho-Nguyen, nam.ho-nguyen@sydney.edu.au

We consider the convex bilevel optimization problem, also known as simple bilevel programming. Convergence of algorithms is in general not guaranteed without further conditions due to the possibility of super-optimal solutions. We present counter-examples to demonstrate this, and explore conditions which guarantee convergence, drawing connections to Lagrange duality and Levitin-Polyak well-posedness. Finally, we discuss the stochastic setting, which is particularly useful in large-scale problems. We develop a stochastic projection-free algorithm that employs variance-reduced gradient estimators to achieve simultaneous dual-level convergence.

An inexact Primal-Dual Newton method for nonconvex optimization with equality constraints

Oscar Smee, o.smee@uq.edu.au

Optimisation problems with non-linear equality constraints are pervasive across science, engineering, and, more recently, machine learning. A prominent class of methods for solving such problems includes Sequential Quadratic Programming (SQP) and, more broadly, primal-dual Newton methods. Traditional analyses of these methods often rely on strong assumptions, most notably that the Hessian of the Lagrangian remains positive definite on the null space of the constraint Jacobian. This requirement typically necessitates strong regularisation or the use of a positive definite proxy for the Hessian. In large-scale settings, solving the primal-dual Newton subproblem exactly is often infeasible, further complicating the issue. In this work, we propose a primal-dual algorithm that incorporates an inner solver based on the conjugate residual method. We leverage several recently discovered properties of conjugate residual, including negative curvature detection, iterate monotonicity, and descent conditions, to develop an approach capable of handling negative curvature within the subproblem as it arises. Our method thereby avoids detrimental regularisation of the Lagrangian Hessian while simultaneously accommodating inexact subproblem solves. Additionally, we establish competitive worst-case global convergence guarantees for our algorithm.

Beyond Discretization: Learning the Optimal Solution Path

Paul Grigas, pgrigas@berkeley.edu

Many applications require minimizing a family of optimization problems, indexed by some hyperparameter vector, to obtain an entire solution path (or solution map). Traditional approaches proceed by discretizing and solving a series of optimization problems. We propose an alternative approach that parametrizes the solution path within a given function class and solves a single stochastic optimization problem to learn the entire solution path. Our method offers substantial complexity improvements over discretization. When using constant-step size SGD, the uniform error of our learned solution path relative to the true path exhibits linear convergence to a constant that diminishes as the expressiveness of the function class increases. In the case of a one-dimensional hyperparameter, we prove stronger results that demonstrate complexity improvements over the best known results for uniform discretization whenever the solution path is sufficiently smooth. Finally, we discuss other extensions including an adaptive variant of our method that sequentially adds basis functions, and we demonstrate strong numerical performance through experiments on imbalanced binary classification and portfolio optimization examples.

Bilinear Logic-Based Benders Decomposition

Ponpot Jartnillaphand, 19877688@student.curtin.edu.au

Bilinear Logic-Based Benders Decomposition (BLBBD) is a new exact framework that strengthens classical Logic-Based Benders Decomposition (LBBD). Standard LBBD often relies on no-good

cuts, which can be weak and slow to converge. BLBBD addresses this by reformulating flattened decision variables into bilinear components, enabling the generation of stronger and more general cuts. The approach retains the LBBD structure, with a master–subproblem framework, but adds an auxiliary bilinear problem that enhances infeasibility analysis and cut derivation. We demonstrate the effectiveness of BLBBD on complex scheduling and routing problems, showing substantial improvements over both commercial solvers and traditional LBBD in terms of scalability, convergence speed, and solution quality.

Multi-Objective Structural Optimization Using Evolutionary Algorithms: Challenges and Future Directions

Saeid Kazemzadeh, saeid.azad@atilim.edu.tr

Multi-objective design optimization of structural systems has long been recognized as an efficient approach to provide reasonable solutions for construction projects. In the past few decades, numerous evolutionary algorithms have been developed and employed for optimization of structural systems considering multiple design objectives. Although evolutionary multi-objective optimization methods can provide alternative no-dominated solutions, they suffer from enormously high number of structural analyses required for structural optimization. Accordingly, computational cost of these algorithms in large scale structural optimization problems is often excessive. In this study, the recent evolutionary multi-objective structural optimization algorithms developed for optimization of steel skeletal structures are outlined. Moreover, contemporary techniques to reduce the total computational time of the structural optimization process are also outlined. Finally, challenges in benchmarking multi-objective structural optimization algorithms are discussed and future research directions are summarized.

Shutdown Maintenance Scheduling using Two-Stage Stochastic Programming with Endogenous Uncertainty

Srimali Gunasekara, c.ekanayak1@postgrad.curtin.edu.au

This talk addresses maintenance scheduling across two consecutive shutdowns. Budget, resource, and time constraints may require postponing maintenance of some components, increasing the risk of costly failures. We minimize total expected costs using a two-stage stochastic programming model: the first stage selects components to postpone, while the second adapts maintenance decisions based on observed failures. To solve this complex problem, we employ a Sample Average Approximation (SAA) framework with scenario sampling. Due to the exponential growth of the decision space, we propose a heuristic supported by solution space reduction techniques, including feasibility screening, elimination of suboptimal solutions, and restricted neighborhood search. Computational experiments on benchmark instances demonstrate the effectiveness of these methods under decision-dependent uncertainty.

Optimality analysis of bilevel programming problems and its application for the electric vehicle charging stations challenges

Tamanna yadav, tamannayadav4@gmail.com

The focus of this talk is on a specific class of optimization problems known as non-smooth bilevel programming problems. Through the utilization of the limiting constraint qualification, a necessary optimality condition for this two-level optimization model is established. Subsequently, using the concept of generalized convexity over cones, sufficient optimality conditions are developed. This approach broadens the understanding of generalized convexity within the context of bilevel programming. Moreover, an application for the electric vehicle charging stations is discussed. To complement the theoretical developments, many numerical illustrations are strategically placed

throughout the presentation to validate the derived results. These illustrations provide concrete examples that demonstrate the practical relevance and effectiveness of the proposed optimization framework.

From Linear to Superlinear: Advances in Derivative-Free Optimisation Algorithms

[Thakshila Rajapaksha](#), thakshila@uow.edu.au

Derivative-free optimisation (DFO) methods are essential when (sub)gradient information is unavailable or unreliable, particularly in nonsmooth and high-dimensional problems. This talk presents new theoretical and numerical insights into the convergence behaviour of two key DFO schemes: a proximal bundle method and the VU-algorithm. First, I establish a linear convergence rate for the DFO proximal bundle framework, showing that, under reasonable error-bound assumptions, both function values and iterates converge at a linear rate. Numerical experiments on piecewise linear-quadratic functions validate the theoretical rates. I then extend the analysis to weakly convex functions, showing how the same framework can be adapted to ensure convergence guarantees beyond the purely convex case. Finally, I provide a brief discussion of the DFO VU-algorithm, a scheme that alternates between steps in “V” and “U” spaces, and highlight its potential to achieve superlinear convergence in nonsmooth optimisation. The analysis indicates that, with sufficiently accurate subgradient and Hessian approximations, the method can accelerate the convergence behaviour of the DFO proximal bundle method, which corresponds to the V-step, by incorporating a quasi-Newton scheme as the U-step. Together, these results advance the theoretical foundations of DFO, point to future research directions centred on the VU framework, and open pathways for practical applications where derivative information is unavailable.

A general approach to distributed operator splitting

[Thang D. Truong](#), thang.tdk64@gmail.com

Splitting methods have emerged as powerful tools to address complex problems by decomposing them into smaller solvable components. In this work, we develop a general approach of forward-backward splitting methods for solving monotone inclusion problems involving both set-valued and single-valued operators, where the latter may lack cocoercivity. Our proposed approach, based on some coefficient matrices, not only encompasses several important existing algorithms but also extends to new ones, offering greater flexibility for different applications. Moreover, by appropriately selecting the coefficient matrices, the resulting algorithms can be implemented in a distributed and decentralized manner.

Circuit realization and hardware linearization of monotone operator equilibrium networks

[Thomas Chaffey](#), thomas.chaffey@sydney.edu.au

This talk will introduce a framework for performing machine learning in analog electronic hardware, grounded in monotone operator theory. First it will be shown that the class of monotone operator equilibrium networks corresponds one-to-one with the class of circuits built from diodes, resistors, transformer and gyrators, allowing inference to be performed in a single operation. It will then be shown that the gradient of such a circuit can be calculated directly in hardware, using a procedure called hardware linearization.

Data-Driven Optimization for Meal Delivery: A Reinforcement Learning Approach for Order-Courier Assignment and Routing at Meituan

[Tomas Lagos](#), tomasignacio.lagos@sydney.edu.au

The rapid growth of online meal delivery has introduced complex logistical challenges, where platforms must dynamically assign orders to couriers while accounting for demand uncertainty, courier

autonomy, and service efficiency. Traditional dispatching methods, often focused on short-term cost minimization, fail to capture the long-term implications of assignment decisions on system-wide performance. This paper presents a novel hybrid framework that integrates reinforcement learning with hyper-heuristic optimization to improve sequential order assignment and routing decisions in meal delivery operations. Our approach combines n -step SARSA with value function approximation and a multi-armed bandit-based hyper-heuristic incorporating seven specialized low-level heuristics. Our approach explicitly models the evolving system state, enabling dispatching policies that balance immediate efficiency with future operational performance. By employing scalable linear value function approximation, we enhance policy learning in high-dimensional environments while maintaining generalization across states and actions. Using real operational data from the food delivery platform Meituan, we develop a comprehensive simulation environment that captures order dynamics, courier behavior, and service times. Through extensive computational experiments, we demonstrate that our framework significantly outperforms traditional benchmark policies, achieving 12% cost reduction through strategic order postponement. Our results reveal that the largest improvements occur during high-demand periods with courier shortages, and that a 10% increase in courier availability yields greater benefits than algorithmic improvements alone. The proposed methodology effectively balances immediate operational efficiency with long-term performance, while providing valuable insights for meal delivery platforms regarding courier fleet management and order assignment strategies.

KKT-based optimality conditions for neural network approximations

[Vinesha Peiris, m.peiris@curtin.edu.au](mailto:m.peiris@curtin.edu.au)

In this talk, we obtain necessary optimality conditions for neural network approximation. We consider neural networks in Manhattan (l_1 norm) and Chebyshev (max norm). The optimality conditions are based on neural networks with at most one hidden layer. We reformulate nonsmooth unconstrained optimisation problems as larger dimension constrained problems with smooth objective functions and constraints. Then we use KKT conditions to develop the necessary conditions and present the optimality conditions in terms of convex analysis and convex sets.