AI-OPT 2023

LOCATION: Deakin Downtown, Tower 2, 727 Collins Street, Melbourne

Day One (Thursday 14th December)

08:30 – 09:00	**Registration & Tea/Coffee**	
09:00 – 09:10	**Welcome and Introduction (Acknowledgement of Country)**	
SESSION ONE		
09:10 – 10:00	**Spotlight Talks**	
	Markus Wagner	ELEA - Build your own Evolutionary Algorithm in your Browser
	Jialiang Li	Evolutionary Computation Enhanced Parameterized Algorithms
	Ishara Hewa Pathiranage	Multi-Objective Evolutionary Algorithms for the Chance-Constrained Knapsack Problem with Normally Distributed Random Variables
	Yunqi (Ryan) Wang	Community Battery Energy Storage Systems for Enhancing Distribution System Operation: A Multi-objective Optimization Approach
	Seyed Mohammad Khalili	Two-Stage Stochastic Programming for Transit-Based Evacuation Planning of Vulnerable People During Natural Hazards
10:00 – 10:40	**Regular Talks**	
	Neelofar Neelofar	Search-based Optimization and Software Engineering
	Milad Dehghan	Exact Mixed-Integer and Constraint Programming Solutions to the Two-Dimensional Bin Packing Problem with Due Dates
10:40 – 11:10	**Morning Tea**	
SESSION TWO		
11:10 – 12:50	Hongjie Xu	Learning-based Ant Colony Pricing Heuristic for Enhancing Column Generation for Bin Packing
	Adam McGregor	A Lookahead Column Generation Based Matheuristic for the Variable-Sized Two-Dimensional Bin Packing Problem
	Miquel Ramirez	Kinodynamic Motion Planning via Branch-and-Cut over Probabilistic Roadmaps
	Gleb Belov	High-Level Modeling for MIP: Status Update
	Rounak Saha Niloy	Multi-Concept Optimization: Challenges and Opportunities
12:50 – 14:00	**Lunch**	

AI-OPT 2023

LOCATION: Deakin Downtown, Tower 2, 727 Collins Street, Melbourne

Day One (Thursday 14th December)

SESSION THREE		
14:00 – 15:40	Denis Antipov	Theory of Diversity Optimization for Multi-Objective Problems
	Kamrul Hasan Rahi	A Generalized Surrogate-Assisted Evolutionary Algorithm for Solving Expensive Multi-Objective Optimization Problems
	Bing Wang	An Evaluation of Simple Solution Transfer Strategies for Bilevel Multi-Objective Optimization
	Adel Nikfarjam	Quality Diversity Approaches for Time Use Optimisation to Improve Health Outcomes
	Mario Andrés Muñoz	A Similarity Measure of Continuous Black-Box Problems
15:40 – 16:10	**Afternoon Tea**	
SESSION FOUR		
16:10 – 17:30	Aneta Neumann	Diversity Optimization for the Detection and Concealment of Spatially Defined Communication Networks
	Rehan Mendis	Optimal Deployment of Smart Meters and Monitoring Devices in Water Networks
	Canchen Jiang	Rolling Horizon Optimization for Network-Aware Value Stacking of Electric Vehicle Coordination Under Uncertainties
	Jiarong Fan	Multi-Agent Reinforcement Learning for Decentralized Electric Vehicle Charging Coordination
17:30 – 18:50	**Networking and Drinks**	

AI-OPT 2023

LOCATION: Deakin Downtown, Tower 2, 727 Collins Street, Melbourne

Day Two (Friday 15th December)

08:30 – 09:00	**Arrival Tea/Coffee**	
SESSION ONE		
09:00 – 10:00	**Spotlight Talks**	
	Markus Wagner	CryptOpt: Verified Compilation with Randomized Program Search for Cryptographic Primitives
	Markus Wagner	Automatic Detection of Fallacies in Climate Change Misinformation
	Zhuo Wei	Fair Solar PV Curtailment in Electricity Distribution Networks Using Reinforcement Learning
	Long Nguyen	Sensor Allocation and Online-Learning-based Path Planning for Maritime Situational Awareness Enhancement: A Multi-Agent Approach
	Qingyang Li	Synthesizing Mixed-Integer Linear Programming Models from Natural Language Descriptions
	Anh-Dzung (Dzung) Doan	Assessing Domain Gap for Continual Domain Adaptation in Object Detection
10:00 – 10:40	**Regular Talks**	
	Sevvandi Kandanaarachchi	Painful Removal of Tiling Artefacts in Hyperspectral Data
	Mingyu Guo	Limited Query Graph Connectivity Test
10:40 – 11:10	**Morning Tea**	
SESSION TWO		
11:10 – 12:50	Jinhao Li	AI-Empowered Real-Time Coordination for Wind and Energy Storage in Australian Wholesale Energy and Ancillary Service Markets
	Amir Fayaz Heidari	Planning for Inertia and Resource Adequacy in a Renewable-rich Power System: A Competitive Market Expansion Approach
	Diksha Goel	Evolving Reinforcement Learning Environment to Minimize Learner's Achievable Reward: An Application on Hardening Active Directory Systems
	Kamal Mammadov	Gradient Descent Learning Algorithms for Finding Stable Fixed Points of Bimatrix Games
	Harry McArthur	Preserving Privacy while Publishing Information
12:50 – 13:50	**Lunch**	

AI-OPT 2023

LOCATION: Deakin Downtown, Tower 2, 727 Collins Street, Melbourne

Day Two (Friday 15th December)

SESSION THREE		
13:50 – 15:30	Xiankun Yan	Optimizing Chance-Constrained Submodular Problems with Variable Uncertainties
	Thilina Chathuranga Pathirage Don	The Chance-Constrained Travelling Thief Problem
	Bahram Farhadinia	Efficient Packing for Coles' Online Orders in Australia with Ant Colony Optimization
	Shizhe Zhao	Battery on Wheel: Optimising Charging/Discharging Schedule to Reduce Emissions and Cost for EV Owners
	Marcus Gallagher	Benchmarking Optimisation Algorithms and Neural Network Loss Landscapes
15:30 – 16:00	**Afternoon Tea**	
SESSION FOUR		
16:00 – 17:20	Angus Kenny	Hybridising Tree-Based Pipeline Optimisation Tool with Bayesian Optimisation
	Yogesh Pipada Sunil Kumar	Constraint Learning Using Relu DNN for an Optimal Bidding Engine for a Flexibility Aggregator in the European Electricity Market
	Pivithuru Thejan	AI-Copilot for Business Optimisation: A Framework and A Case Study in Production Scheduling
	Edward Lam	Learning to Branch in Integer Programming
17:20 – 17:30	**Wrap up**	

The 2023 Workshop on AI-based Optimisation

AI-OPT 2023

December 14-15, 2023

Deakin Downtown, Tower 2, 727 Collins Street, Melbourne

BOOK OF ABSTRACTS

Spotlight Talks

Day 1, 09:10 – 10:00

ELEA - Build your own Evolutionary Algorithm in your Browser

Markus Wagner

We provide an open-source framework to experiment with evolutionary algorithms which we call "Experimenting and Learning toolkit for Evolutionary Algorithms (ELEA)". ELEA is browser-based and allows to assemble evolutionary algorithms using drag-and-drop, starting from a number of simple pre-designed examples, making the startup costs for employing the toolkit minimal. The designed examples can be executed and collected data can be displayed graphically. Further features include export of algorithm designs and experimental results as well as multi-threading. With the very intuitive user interface and the short time to get initial experiments going, this tool is especially suitable for explorative analyses of algorithms as well as for the use in classrooms.

URL: https://elea-toolkit.netlify.app/

Evolutionary Computation Enhanced Parameterized Algorithms

Jialiang Li

NP-hard problems have fundamental influence across various industries. Unfortunately, it is widely believed intractable to exactly solve any of these problems. To neutralize the intractability, parameterized complexity and associated algorithms have been developed with the purpose of understanding the implicit structures of NP-hard problems. These structures are generally interpreted by the combinations of different parameters. The typical choices are including the cardinalities of the solutions or hand-crafted parameters like tree-width. Tree-width is a prominent graph parameter employed to measure the resemblance between an arbitrary graph and a tree (or a forest). Many NP hard graph combinatorial problems become 'tractable' when the tree-width falling into a preferable range. Therefore, it is rational to believe that if the tree-width stays within a threshold, other exact and approximation algorithmic paradigms, such as dynamic programming, would also benefit. However, for general graphs, it is unrealistic to expect that the tree-widths will always remain within a desired threshold. In this work, we investigate the potential of evolutionary computations on accelerating dynamic programming on tree decompositions by strategically reduce the tree-width. We proposed a zoom-in evolutionary algorithm and use empirical experiments to demonstrate the generalizability of this method.

Multi-Objective Evolutionary Algorithms for the Chance-Constrained Knapsack Problem with Normally Distributed Random Variables

Ishara Hewa Pathiranage

Optimizing real-world problems often involves decision-making under uncertainty due to the presence of unknown or uncontrollable variables. In such cases, the optimal solution may not be feasible or lead to undesirable outcomes. Therefore, it becomes necessary to consider the stochastic components when modeling optimization problems. A chance constraint is a constraint that can be violated with a small probability. Chance-constrained optimization problems allow for modeling such problems by ensuring the meeting of the constraint with high confidence. Evolutionary algorithms are bio-inspired algorithms successfully employed in various contexts to achieve high-quality results, including chance-constrained optimization problems. We investigate the use of multi-objective evolutionary algorithms (MOEAs) for solving the stochastic knapsack problem, where the weights are stochastic which are independent and normally distributed. We introduce a three-objective approach to obtain a set of solutions that includes an optimal solution for any possible confidence level imposed on the constraint. We compare this approach with a two-objective formulation, which works only with one confidence level. We evaluate the problem using two different MOEAs (GSEMO and MOEA/D variants) across various benchmark scenarios. Overall, we demonstrate that both formulations have their own advantages, and the three-objective approach performs similarly to the two-objective approach, indicating that the cost of universality for this approach is minimal.

Community Battery Energy Storage Systems for Enhancing Distribution System Operation: A Multi-objective Optimization Approach

Yunqi (Ryan) Wang

The growing penetration of distributed energy resources (DERs) in distribution networks (DNs) raises new operational challenges, particularly in terms of reliability and voltage regulation. In response to these challenges, we introduce an innovative DN operation framework with multi-objective optimization, leveraging community battery energy storage systems (C-BESS). The proposed framework targets two key operational objectives: first, to minimize voltage deviation, which is a concern for a distribution network service provider (DNSP), and second, to maximize the utilization of DERs on the demand side. Recognizing the conflicting nature of these objectives, we utilize C-BESS to enhance the system's adaptability to dynamically adjust DN operations. The multi-objective optimization problem is solved using the non-dominated sorting genetic algorithm-II (NSGA-II). Case studies using real-world data are conducted to validate the effectiveness of the proposed framework. The results show significant improvements in voltage regulation and DER utilization, demonstrating the potential of C-BESS in enabling more reliable DN operation. Our findings contribute to the ongoing discourse on the role of C-BESS in DN operation enhancement and DER integration.

Two-Stage Stochastic Programming for Transit-Based Evacuation Planning of Vulnerable People During Natural Hazards

Seyed Mohammad Khalili

The impact of natural hazards is escalating globally, affecting diverse regions. Community resilience depends significantly on adaptive disaster management strategies and preparedness measures. Transit-based evacuation planning is a vital component of disaster management, focused on transporting transit-dependent evacuees, particularly vulnerable individuals with limited private transportation access, from affected areas to safe shelters within strict time constraints crucial for saving lives. These plans can be formulated through mathematical modelling approaches, sharing similarities with vehicle routing problems. This study proposes a two-stage scenario-based mixed-integer stochastic programming mathematical model. In the first stage, the model determines the optimal location and required number of vehicle depots in disaster-prone areas. In contrast, in the second stage, it optimizes vehicle routing and scheduling in various disaster scenarios. These scenarios encompass different predicted disaster impacts on the transportation network and their likelihoods, involving factors such as road closures, rendering some parts of the network inaccessible and significantly impacting travel times. Although the likelihood of these disaster scenarios is typically low, their impact can be devastating, constituting disruption risks. Consequently, we formulated the objective function of this problem using the total evacuation time via the Conditional Value at Risk (CVaR) measure, renowned for its ability to quantify disruption risks. The proposed model and its applicability were validated through a real-life case study in the flood-prone areas of northern New South Wales.

Day 2, 09:00 – 10:00

CryptOpt: Verified Compilation with Randomized Program Search for Cryptographic Primitives#

Markus Wagner

Most software domains rely on compilers to translate high-level code to multiple different machine languages, with performance not too much worse than what developers would have the patience to write directly in assembly language. However, cryptography has been an exception, where many performance-critical routines have been written directly in assembly (sometimes through metaprogramming layers). Some past work has shown how to do formal verification of that assembly, and other work has shown how to generate C code automatically along with formal proof, but with consequent performance penalties vs. the best-known assembly. We present CryptOpt, the first compilation pipeline that specializes high-level cryptographic functional programs into assembly code significantly faster than what GCC or Clang produce, with mechanized proof (in Coq) whose final theorem statement mentions little beyond the input functional program and the operational semantics of x86-64 assembly. On the optimization side, we apply randomized search through the space of assembly programs, with repeated automatic benchmarking on target CPUs. On the formal-verification side, we connect to the Fiat Cryptography framework (which translates functional programs into C-like IR code) and extend it with a new formally verified program-equivalence checker, incorporating a modest subset of known features of SMT solvers and symbolic-execution engines. The overall prototype is quite practical, e.g., producing new fastest-known implementations of finite-field arithmetic for both Curve25519 (part of the TLS standard) and the Bitcoin elliptic curve secp256k1 for the Intel 12th and 13th generations.

Automatic Detection of Fallacies in Climate Change Misinformation

Markus Wagner

Misinformation about climate change causes a number of negative societal effects, necessitating interventions to reduce its influence. Technique-based corrections—exposing the logical fallacies and misleading rhetorical techniques in misinformation—has been shown to be an effective approach to reducing the influence of misinformation. However, little progress has been made on documenting and detecting fallacies in climate misinformation. In this study, wet apply a previously developed critical thinking methodology for deconstructing climate misinformation, in order to develop a dataset mapping different types of climate misinformation to reasoning fallacies. This dataset is used to train a model to detect fallacies in climate misinformation. The fallacies that are easiest to detect include fake experts, ad hominem attacks, and anecdotal arguments. Fallacies that require background knowledge, such as oversimplification, misrepresentation, and slothful induction, are relatively more difficult to detect. Fallacies that conceptually overlap with other fallacies such as conspiracy theories and ad hominem are also a source of difficulty in detection. This research lays the groundwork for development of solutions where automatically detected climate misinformation can be countered with generative technique-based corrections.

Fair Solar PV Curtailment in Electricity Distribution Networks Using Reinforcement Learning

Zhuo Wei

The widespread adoption of residential solar photovoltaics (PV) has shown great potential. However, it has also introduced a set of challenges, one of which is the occurrence of frequent overvoltage events. These overvoltage problems are the results of reverse power flows, wherein excess energy generated by residential PV systems flows back into the grid. To minimize the risk of damage to devices connected to the distribution network, PV inverters typically employ a curtailment strategy. This strategy reduces the power fed into the grid by reducing PV power export, thus mitigating the voltage rise.

While PV curtailment is a necessary measure, it can also lead to unintended results. Households at the far end of the feeder (i.e., power line in the low-voltage electric distribution network) will experience more frequent and substantial PV curtailment. This disparity arises because overvoltage is unevenly distributed across the distribution network, resulting in higher voltages at households nearer to the end of the feeder. As a result, households located near the end of the feeder receive lower value from the same-sized PV panels, leading to an inequitable allocation of the benefits of PV generation. Therefore, the curtailment of PV energy needs to be optimized to ensure both safety and fairness. However, achieving this optimization can be difficult because it requires precise information on the distribution network's parameters, which are often unknown or challenging to obtain in real-world scenarios. This challenge motivates us to explore model-free approaches to enable fair PV curtailment.

The focus of our work is to investigate Reinforcement Learning (RL) as a model-free solution to the fair PV curtailment problem. RL is a machine learning technique that can gradually optimize a fair PV curtailment strategy by interacting with the distribution network without knowing the explicit form of the network model. To evaluate the effectiveness of the RL algorithm in achieving fair PV curtailment, we consider six different fairness definitions. Our results demonstrate that all the defined fairness cases permit effective learning of a control policy for curtailing PV export in a fair manner. The findings suggest that RL holds significant promise as a model-free approach to ensure both overvoltage safety and fairness, making it a suitable replacement for existing curtailment strategies.

Sensor Allocation and Online-Learning-based Path Planning for Maritime Situational Awareness Enhancement: A Multi-Agent Approach

Long Nguyen

Countries with access to large bodies of water often aim to protect their maritime transport by employing maritime surveillance systems. However, the number of available sensors (e.g., cameras) is typically small compared to the to-be-monitored targets, and their Field of View (FOV) and range are often limited. This makes improving the situational awareness of maritime transports challenging. To this end, we propose a method that not only distributes multiple sensors but also plans paths for them to observe multiple targets, while minimizing the time needed to achieve situational awareness. In particular, we provide a formulation of this sensor allocation and path planning problem which considers the partial awareness of the targets' state, as well as the unawareness of the targets' trajectories. To solve the problem, we present two algorithms: 1) a greedy algorithm for assigning sensors to targets, and 2) a distributed multi-agent path planning algorithm based on regret-matching learning. Because a quick convergence is a requirement for algorithms developed for high mobility environments, we employ a forgetting factor to quickly converge to correlated equilibrium solutions. Experimental results show that our combined approach achieves situational awareness more quickly than related work.

Synthesizing Mixed-Integer Linear Programming Models from Natural Language Descriptions

Qingyang Li

Numerous real-life problems arising in various domains such as production planning, resource allocation, scheduling, and transportation planning and management, can be effectively formulated and solved using Mixed-Integer Linear Programming (MILP) models. However, the transformation of real-world decision-making problems into MILP models heavily relies on expertise in operations research and mathematical optimization, which restricts non-experts' accessibility to MILP. To address this challenge, we propose a framework for automatically formulating MILP models from unstructured natural language descriptions of decision problems, which integrates Large Language Models (LLMs) and mathematical modeling techniques. This framework consists of three phases: i) identification of decision variables, ii) classification of objective and constraints, and iii) finally, generation of MILP models.

In this study, we propose a knowledge representation structure that includes relational knowledge, inheritable knowledge, and so on. Specifically, we present a constraint classification scheme and a set of constraint templates that can guide the LLMs in synthesizing a complete MILP model from a given natural language description of a decision-making problem. After fine-tuning LLMs based on numerous examples of constraint and objective descriptions, our approach can, from the natural language description, identify and synthesize logic constraints in addition to classic demand and resource constraints. The logic constraints have not been studied in existing work.

To evaluate the performance of the proposed framework, we extend the NL4OPT dataset with more problem descriptions and constraint types, in particular logic constraints and binary variables, and with the newly created dataset, we compare our method with one-step model generation methods offered by LLMs like ChatGPT and Google Bard. The experimental results reveal that with respect to the model-generation accuracy and the constraint-generation accuracy, our method which integrates knowledge representation with LLMs significantly outperforms the others. Taking ChatGPT for example, our framework accurately formulates 86.67% problems and 99.32% constraints compared to ChatGPT's 26.67% and 82.99% respectively. The three-phase framework that we developed is a prototype system for synthesizing MILP models from natural language descriptions and has a great potential to capture more constraints for more complex MILPs. Our framework opens up opportunities for developing training tools for operations research practitioners and has the potential to be a powerful tool for automatic decision problem modeling and solving in practice.

Assessing Domain Gap for Continual Domain Adaptation in Object Detection

Anh-Dzung (Dzung) Doan

To ensure reliable object detection in autonomous systems, the detector must be able to adapt to changes in appearance caused by environmental factors such as time of day, weather, and seasons. Continually adapting the detector to incorporate these changes is a promising solution, but it can be computationally costly. Our proposed approach is to selectively adapt the detector only when necessary, using new data that does not have the same distribution as the current training data. To this end, we investigate three popular metrics for domain gap evaluation and find that there is a correlation between the domain gap and detection accuracy. Therefore, we apply the domain gap as a criterion to decide when to adapt the detector. Our experiments show that our approach has the potential to improve the efficiency of the detector’s operation in real-world scenarios, where environmental conditions change in a cyclical manner, without sacrificing the overall performance of the detector. Our code is publicly available https://github.com/dadung/DGE-CDA.

Regular Talks

Day 1, Session One, 10:00 – 10:40

Search-based Optimization and Software Engineering

Neelofar Neelofar

Search-based software testing (SBST) has evolved into a mature field, offering various techniques to address the complex challenges associated with software testing. While these techniques have demonstrated promising outcomes for conventional software and have found successful applications in industry for automated test case generation, their effectiveness is contingent on the specific problem at hand. Consequently, it is crucial to identify problem-specific features that elucidate why a particular instance poses difficulties for an SBST technique, thereby discerning the strengths and weaknesses of these approaches.

For AI-based systems, such as Autonomous Vehicles (AVs), ensuring safety is paramount. Simulated environments offer a safer alternative for testing than in-field operational tests. However, creating a comprehensive test suite to identify critical scenarios is computationally demanding, given the intricate representation of each test, encompassing dynamic and static features such as the AV under test, road participants (vehicles, pedestrians, and static obstacles), environmental factors (weather and lighting), and the road's structural features (lanes, turns, road speed, etc.). Therefore, there is a need for standardized frameworks to alleviate the "curse of dimensionality," making it feasible to apply search-based optimization methods to solve such problems.

In this presentation, I will delve into my work related to search-based software testing, addressing challenges in both traditional and AI-based software systems.

Exact Mixed-Integer and Constraint Programming Solutions to the Two-Dimensional Bin Packing Problem with Due Dates

Milad Dehghan

The two-dimensional non-oriented bin packing problem with due dates consists in packing a set of rectangular items, which may be rotated by 90 degrees, into identical rectangular bins. The bins have equal processing times. The problem searches for a feasible layout of items so as to minimize the maximum lateness, with the lateness of each item being the difference between its due date and the completion time of its assigned bin. This minimax problem is computationally challenging: It has a huge solution space with many alternative solutions whose packing configurations are either symmetric or lead to the same solution cost.

We propose two exact approaches. First, we solve the problem as a mixed-integer program (MIP), which we strengthen with a set of feasibility and symmetry-breaking constraints. Our second solution employs constraint programming (CP), and thus benefits from the strength of cumulative scheduling relaxations of the packing problem. Our extensive computational experiments show that MIP, which stands as a more traditional modelling technique within the cutting and packing community, is only successful on small-sized benchmark instances with as few as 20 items. It is outperformed by CP that can solve medium and large instances with up to 100 items in a reasonable time.

Learning-based Ant Colony Pricing Heuristic for Enhancing Column Generation for Bin Packing

Hongjie Xu

Column generation is a technique used in combinatorial optimization to solve large-scale problems more efficiently. The core idea behind column generation is to start with a subset of variables (columns) and iteratively add new columns that improve the objective function. The process often involves solving a pricing problem, which identifies the most promising column to add to the model. However, during CG procedure, the efficiency of solving pricing problem heavily determines the overall performance as pricing problems are difficult to solve. Multiple pricing is involved to generate multiple columns in a heuristic way, instead of generating only one column in each iteration. Effectively generating multiple high-quality columns is a crucial research topic in Column Generation for solving large-scale combinatorial optimization problems. Ant Colony Optimization (ACO) algorithms, a population-based method, are a class of probabilistic techniques known for their versatility in sampling multiple high-quality solutions. However, effectiveness of ACO algorithms highly depends on the tuning of various parameters as these parameters control how the algorithm behaves during the searching process. In this paper, we introduced a Machine learning based method that combines machine learning and ACO to effectively sample high quality columns. Specifically, we extracted features from solved small instances during column generation procedures. These features are subsequently employed to train a classification model that can predict the likelihood of a decision variable being chosen for inclusion in an optimal column. We then integrated our classification model with ACO by utilizing the predicted probabilities of potential decision variables. This integration significantly improves ACO's column sampling process. We evaluated our proposed method on a well-known combinatorial optimization problem, Bin packing problem (BPP). Our experimental results show that our method could significantly improve the quality of columns generated using ACO as compared to other heuristic and exact pricing problem solvers. In addition, we extend our proposed method to enhance Branch-and-price as the pricing problem could be solved efficiently.

A Lookahead Column Generation Based Matheuristic for the Variable-Sized Two-Dimensional Bin Packing Problem

Adam McGregor

Given a set of small rectangular items and large non-identical rectangular bins, the variable-sized two-dimensional guillotine bin packing problem searches for a feasible packing of the items into the bins, with the items obtained via guillotine cuts. It minimizes the total cost of the used bins. We propose a hybrid column generation based matheuristic that solves the problem approximately. Its reduced master problem is the classical Dantzig-Wolfe formulation of the problem, which is solved as an integer program. The master problem relies on a collection of feasible packings that are constructed heuristically during the search process. The initial set of packings is obtained via a constructive heuristic, which selects and packs bins sequentially. Solving a series of mixed-integer programs, it implements a look-ahead mechanism that prohibits the investigation of infeasible directions and constrains the search to improving ones. The constructive heuristic stops when it exceeds the allocated time. The metaheuristic then attempts to enhance the incumbent solution. It solves a sequence of pricing problems examining each of the bin types. It searches for a feasible packing with the largest negative reduced cost, which it adds as a column to the master problem. The search terminates either when the time limit is exceeded or when the approach fails to find a solution to the pricing problem. Packing of a bin is always done approximately via an algorithm that hybridizes constraint programming with a heuristic search. When compared to the state of the art, the extensive computational experiments provide evidence of the good performance of the proposed matheuristic.

Kinodynamic Motion Planning via Branch-and-Cut over Probabilistic Roadmaps

Miquel Ramirez

We revise the classic Lazy Probabilistic Roadmaps algorithm (Lazy PRM) for robot motion planning by pairing PRM with a novel Branch-and-Cut (BC) algorithm. Cuts are dynamically generated constraints that are imposed on minimum cost paths over the geometric graphs selected by PRM. Cuts eliminate paths that cannot be mapped into smooth plans that satisfy suitably defined geometric (collision avoidance) and differential (velocity limits) constraints. We generate candidate smooth plans by fitting Hermitian splines to the vertices selected by minimum-cost paths. Constraint satisfaction is then determined with a novel, recently proposed algorithm that maps smooth plans into a finite trace, without the need to choose a fixed discretization step. A trace records the exact sequence of constraint boundaries crossed by the plan, modulo arithmetic precision, and enables deciding whether these crossings amount to a constraint being violated. Both plans and constraint boundaries are described by polynomial functions, a notoriously difficult problem of wide interest in the verification of control systems and mathematical programming. The constraint verification procedure thus can be reused to be integrated with existing optimization algorithms. We evaluate several planners using our proposed methods over the recently proposed BARN benchmark for robot motion planning, reporting evidence of the feasibility and scalability of our approach.

High-Level Modeling for MIP: Status Update

Gleb Belov

We review current state of Constraint Programming-style modeling for Mixed-Integer Programming and automatic reformulation systems. In particular we ask the question whether such modeling is well-defined.

Multi-Concept Optimization: Challenges and Opportunities

Rounak Saha Niloy

Multi-concept optimization (MCO) is a new and powerful approach to design optimization that allows engineers and scientists to concurrently identify the best concept and the corresponding variable values to optimize certain objective(s). This is in contrast to traditional design optimization approaches, which typically involve optimizing a single concept at a time. MCO is particularly well-suited for problems where concept selection is an inherent part of design process. For example, consider the optimization of a cantilever beam with objectives to minimize weight and deflection at the free end. Instead of restricting the search to a pre-determined shape of the cross-section, one may wish to start off by considering beams with multiple different cross-sections (e.g. circular, rectangular, L-shaped, I-shaped or T-shaped cross-sections) as plausible candidate designs, each of which represents a different concept. In fact, some studies estimate that the decisions made during the early phases of design, including concept selection and preliminary design, impact up to 70% of the overall product life-cycle costs. While such problems are commonly encountered in practical domains such as engineering, transport, product design, there has been little focus on developing computationally efficient algorithms for MCO. This is a significant barrier to the adoption of MCO in practice. MCO has a number of advantages over traditional design approaches. First, it can lead to better design solutions by exploring a wider range of concepts and variable values. Second, it can be more efficient, as it can eliminate the need to independently and extensively iterate between different concepts. However, there are a number of challenges in developing computationally efficient algorithms for MCO. This talk will discuss the challenges and opportunities in developing computationally efficient algorithms for MCO. The talk will also present a number of case studies where MCO has been used to solve real-world design problems. This talk is intended for engineers, scientists, and other professionals who are interested in learning more about MCO and its potential applications.

Day 1, Session Three, 14:00 – 15:40

Theory of Diversity Optimization for Multi-Objective Problems

Denis Antipov

The diversity optimization aims at finding a set of solutions with a maximum diversity (according to some measure), all of which satisfy some quality threshold. Despite it has become a hot topic in the optimization community, the theoretical knowledge of how the diversity optimization algorithms work is very limited. Obtaining such result might give us a better understanding of the most effective ways to optimize the diversity and therefore to improve the existing algorithms.

In this talk we study the diversity optimization in the multi-objective setting, where the main goal is to find the most diverse set of solutions covering the whole Pareto-front. We study a simple multi-objective evolutionary algorithm GSEMO on benchmark problems OneMinMax and LeadingOnes-TrailingZeros. For both problems we show that implementing a simple tie-breaking rule which optimizes the diversity into GSEMO allows us to optimize the diversity at the same time, which we need to find the Pareto-front. At the same time, it does not slow down the process of covering the whole Pareto-front. This improves the previously known runtime bounds for OneMinMax and demonstrates that the diversity optimization is not much more complex than the regular multi-objective optimization.

A Generalized Surrogate-Assisted Evolutionary Algorithm for Solving Expensive Multi-Objective Optimization Problems

Kamrul Hasan Rahi

Expensive multi-objective optimization problems (EMOPs) refer to those wherein evaluation of each candidate solution incurs a significant cost. To solve such problems within a limited number of solution evaluations, surrogate-assisted evolutionary algorithms (SAEAs) are often used. However, existing SAEAs typically operate in a generational framework wherein multiple solutions are identified for evaluation in each generation. There exist relatively few proposals in steady-state framework, wherein only a single solution is evaluated in each iteration. The development of such algorithms is crucial to efficiently solve EMOPs for which the evaluation of candidate designs cannot be parallelized. Furthermore, regardless of the framework used, the performance of current SAEAs tends to degrade when the Pareto front (PF) of the problem has irregularities, such as extremely concave/convex segments, even for 2/3-objective problems. To contextualize the motivation of this study, the performance of a few state-of-the-art SAEAs is first demonstrated on some such selected problems. Then, to address the above research gaps, we propose a generalized surrogate-assisted EA (GSAEA) which is able to seamlessly switch between generational and steady-state forms based on the application at hand by simply prescribing number of true evaluations per generation. GSAEA mainly incorporates a number of novel elements including (a) effective use of model uncertainty information to aid the search, including the use of probabilistic dominance and Mahalanobis distance, (b) two-step infill identification using non-dominance (ND) and distance-based selection, and (c) a shadow ND mechanism to avoid repeated selection and evaluation of dominated solutions. The efficacy of the proposed approach is demonstrated through extensive benchmarking on a range of test problems. It shows competitive performance relative to many state-of-the-art SAEAs, including both steady-state and generational approaches. GSAEA has also been used in design optimization of a Vertical-Axis Wind Turbine (VAWT) which involved physical experiments in-loop. This study strengthens the case for applying SAEAs for design optimization in general.

An Evaluation of Simple Solution Transfer Strategies for Bilevel Multi-Objective Optimization

Bing Wang

Bilevel optimization problem (BLOP) refers to a class of problems with a hierarchical structure, wherein a lower-level optimization problem acts as a constraint for an upper-level optimization problem. Evolutionary algorithms (EAs) have been commonly used to solve BLOPs where underlying functions are black-box or do not conform to certain mathematical properties. One of the downsides of using EAs, especially in a nested format, is the significant computational expenses (number of function evaluations); and a number of strategies have been proposed to mitigate this. However, most of the existing studies in the domain of evolutionary bilevel optimization are directed towards problems with single objective at both levels, while very few have explored BLOPs with multiple objectives at one or both levels~(BLMOPs). In this study, we investigate the potential benefits of utilizing knowledge transfer by seeding initial population from neighboring solutions for solving BLMOPs. Towards this end, we construct two simple strategies, referred to as full population transfer and selective population transfer, and study their potential to improve the performance over the baseline nested EA for BLMOPs. Experimental results show that the selective transfer strategy has more reliable and competitive performance compared to baseline. Empirical analysis is presented to highlight the relevant factors that lead to the observed performance trends.

Quality Diversity Approaches for Time Use Optimisation to Improve Health Outcomes

Adel Nikfarjam

Several studies have established a link between people's well-being and how they spend their time. As a result, the problem of how best to allocate finite and constant sum-constrained time, called compositional data, has recently gained significant attention. In other words, compositional data relates to how individuals allocate their time among different daily activities. Here, we utilise a number of (objective) functions derived using compositional data analysis techniques and a large child cohort (n > 1000), to predict how time allocation is associated with health outcomes such as body mass index, life satisfaction, and cognition.

In addition to compositional solution space constraints requiring exhaustive partitioning of a 24-hour day into time-use categories, we also develop and advocate joint cumulative distribution function constraints. Such constraints not only ensure that the feasible solutions do not extrapolate the sampled data for which the objective function is derived from, but ensures potential solutions are realistic/achievable for children in similar populations. Moreover, we incorporate quality diversity approaches to study these objective functions. First, we define a 3-D behavioural space based on the types of activities and aim to find the best-performing time allocations for each objective. This behavioural space is similar to the objective function's search space. Second, we define 2-D behavioural spaces on a pair of objectives and seek to explore the space to find the best time allocations in terms of the third objective in different regions of the spaces.

Finally, we demonstrate a web application, Time allocation optimiser, providing personalised and dynamic, optimised time-use allocation plans where users can select their own desirable range (min/max) for each of the daily activities, producing a visualisation of the corresponding objective function in the feasible solution space.

A Similarity Measure of Continuous Black-Box Problems

Mario Andrés Muñoz

Essential steps in automated algorithm selection are to select a representative set of test instances, and determine an informative representation of them, such that a predictive model can choose the algorithm likely to perform best. Unfortunately, these two steps can strongly influence each other, since a "good" representation often derives from the choice of test instances and algorithms. Since, a priori, we cannot be sure that a constructed benchmark set represents the space of problems, then we should suspect of any representation that makes too many assumptions, as it may present bias. In this talk, I discuss a similarity measure of continuous black-box optimisation problems that uses as few assumptions about the structure of the problem as possible. Constructed as the distance between the joint probability distributions of the input and output variables, the measure can be calculated between problems of the same dimensionality. We compare its performance on some tasks with several baselines, including a more "structured" approach, Exploratory Landscape Analysis, and discuss the advantages and disadvantages of the approach.

Day 1, Session Four, 16:10 – 17:30

Diversity Optimization for the Detection and Concealment of Spatially Defined Communication Networks

Aneta Neumann

In recent years, computing diverse sets of high-quality solutions for an optimization problem has become an important topic. The goal of computing diverse sets of high-quality solutions is to provide a variety of options to decision makers, allowing them to choose the best solution for their particular problem. In this talk, we consider the problem of constructing a wireless communication network for a given set of entities. Our goal is to minimize the area covered by the senders' transmissions while also avoiding adversaries that may observe the communication. We provide evolutionary diversity optimization (EDO) algorithms for this problem. We provide a formulation based on minimum spanning forests that are used as a representation and show how this formulation can be turned into a wireless communication network that avoids a given set of adversaries. We evaluate our EDO approach based on a number of benchmark instances and compare the diversity of the obtained populations in respect to the quality criterion of the given solutions as well as the chosen algorithm parameters. Our results demonstrate the effectiveness of our EDO approaches for the detection and concealment of communication networks both in terms of the quality and the diversity of the obtained solutions.

Optimal Deployment of Smart Meters and Monitoring Devices in Water Networks

Rehan Mendis

Water utility providers supply water to consumers from an underground pipe network. Some of these pipes will develop leaks over time, which can remain undetected for a long time, causing further damage and resulting in significant losses of water from the network.

South East Water (SEW), a water utility in metropolitan Melbourne, has developed a vibration sensor device named Sotto that can detect leaks within the water network. Since the deployment and monitoring of large numbers of sensors has a significant cost, SEW needs to determine how many sensors are required and where they should be placed to detect the majority of leaks.

Our research focuses on developing an optimisation algorithm minimises the overall cost of sensors placed on the one hand, and water lost through undetected leaks on the other hand. This optimisation involves statistical modelling of sensor sensitivity and leak occurrence, both of which vary with characteristics like pipe material and age within the network.

Methodology:

1. Statistical modelling. Our Sensor Sensitivity Model needs to predict the behaviour of sensors with respect to given leak sizes and pipe characteristics along the path from a potential leak location to a potential sensor location. Three binary prediction models were fitted (Lasso, Ridge and Logistic), where a sensor detecting a leak is a positive response, and a sensor not detecting is a negative response. The logistic model produced more accurate predications compared to Lasso and Ridge models.

Our Leak Occurrence Model needs to predict expected leak sizes and probabilities for a given location in a network based on characteristics like pipe material and pipe age. In a preliminary study we have developed a statistical model that can predict leak occurrences in the network, and we are planning to extend this model to also predict leak size.

2. Optimisation. Our aim is to determine the optimal number and placement of sensors in the network based on the statistical models for sensor sensitivity and leak occurrence. We have created a mixed-integer linear formulation of this problem, minimising the overall expected cost of placing sensors and undetected leaks. The statistical models have been converted into constraint functions and implemented using piecewise linear approximations. Based on initial results of the study, the optimisation model shows potential to minimise the loss through undetected water leaks, offset by the cost of the deployed sensors.

Rolling Horizon Optimization for Network-Aware Value Stacking of Electric Vehicle Coordination Under Uncertainties

Canchen Jiang

The global surge in popularity of Electric Vehicles (EVs) is being driven by advancements in battery technology and an unwavering commitment to a zero-emission transport future. This widespread adoption has catalyzed the development of Vehicle-to-Anything (V2X) technologies, which include Vehicle-to-Home (V2H), Vehicle-to-Grid (V2G), and energy trading of EVs in the local grid. While residential EV coordination offers grid-enhancing potential and benefits for EVs, the complexities of V2X operations, involving charging and discharging, pose challenges such as voltage fluctuations in the distribution network. Additionally, the performance of V2X is influenced by uncertainties within the energy system, including renewable energy generation and household energy consumption.

Our work develops a rolling-horizon optimization (RHO) problem for V2X value stacking to fully unlock the value of EV coordination, considering power network constraints (such as voltage limits) and uncertainties including energy consumption and renewable generation in the energy system. By coordinating EVs to perform V2H, V2G, and energy trading, our approach exploits the most valuable services in real-time. We also analyze the expected extra costs caused by the prediction errors to evaluate the impact of uncertainties on the performance of V2X value stacking. To validate our value-stacking model, our work also considers two retail tariffs—time-of-use (TOU) pricing and two-part tariff (TPT) pricing and uses real data from three energy markets, including Australia's National Electricity Market (NEM), ISO New England (ISO-NE), and New York ISO (NY-ISO) in the US.

The numerical results show that V2X value stacking achieves significant benefits to EVs in terms of energy cost reduction. The assessment of the marginal contributions of three value streams shows that energy trading contributed the most to cost reduction in all three energy markets. The uncertainty in energy consumption has a higher impact on the value-stacking performance than PV generation. This highlights the critical need for accurate energy consumption predictions to mitigate the effects of prediction errors on EV value-stacking.

Multi-Agent Reinforcement Learning for Decentralized Electric Vehicle Charging Coordination

Jiarong Fan

The transition towards sustainability in the transport sector underscores the importance of transportation electrification. As the adoption of electric vehicles (EVs) continues to grow, the efficient utilization of EV charging infrastructure and smart charging management emerge as a crucial challenge that calls for innovative solutions.

Different from existing studies, our work focuses on the integration of vehicle-to-vehicle (V2V) energy exchange, enhancing the flexibility of EV charging coordination. This enables EVs with more energy charged to help other EVs that need energy when the grid is stressed. Recognizing the importance of user experience, our work also takes into account the metrics, such as charging satisfaction and fairness. Nonetheless, the complex EV charging environment, including the uncertainties in EV charging behavior and other operational factors, adds complexity to the challenges addressed in our work.

To address the challenges in EV charging coordination with V2V energy exchange, our work presents a new approach based on Multi-Agent Reinforcement Learning (MARL). Our strategy is designed to handle the intricacies of EV charging coordination, particularly the uncertainties associated with EV arrival time, energy prices, and solar energy generation. To enhance the exploration capabilities, our work incorporates an innovative technique: the infusion of parameter noise within the neural network models of MARL. This enhancement accelerates the convergence of the MARL algorithm. Moreover, our work employs the centralized training and decentralized execution technique to enhance the scalability of the strategy while getting access to global information during training to enhance the effectiveness of the strategy.

Our numerical experiments validate the effectiveness of the proposed method. When compared to traditional optimization benchmarks, our MARL-based approach exhibits superior performance under various operational uncertainties and also in terms of scalability as the number of EVs becomes larger. A noteworthy feature of our approach is its decentralized execution, which enhances system reliability, allowing it to effectively manage partial system faults in charging stations.

Day 2, Session One, 10:00 – 10:40

Painful Removal of Tiling Artefacts in Hyperspectral Data

Sevvandi Kandanaarachchi

Often scientists want to get more information about solid samples such as tumour tissues or polymer arrays. TOF-SIMS is an imaging mass spectrometry technique that is used for this task. It works by rastering a pulsed beam of ions across the area of interest. This process results in a visually appealing 3-dimenisonal data cube. Insights on structural and molecular elements can be gained by analysing this data cube.

However, there is a common problem when samples are big (e.g., 20mm x 20mm). A single take of the rastering only covers a small tile of the sample. As such, the sample needs to be moved and the rastering repeated multiple times to image the entire sample. This causes tiling artefacts in the data. That is, we see unwanted criss-cross patterns at tile boundaries. These artefacts are often messy. Sometimes, certain regions of the tile are illuminated more than others.

In this talk, we will explore statistical methods that can help remove tiling artefacts in TOF-SIMS data. We will compare tensor decomposition, discriminant analysis, linear models and other techniques and discuss our findings.

Limited Query Graph Connectivity Test

Mingyu Guo

We propose a combinatorial optimisation model called Limited Query Graph Connectivity Test. We consider a graph whose edges have two possible states (On/Off). The edges' states are hidden initially. We could query an edge to reveal its state. Given a source s and a destination t, we aim to test s-t connectivity by identifying either a path (consisting of only On edges) or a cut (consisting of only Off edges). We are limited to B queries, after which we stop regardless of whether graph connectivity is established. We aim to design a query policy that minimizes the expected number of queries.

Our model is mainly motivated by a cyber security use case where we need to establish whether an attack path exists in a network, between a source and a destination. Edge query is resolved by manual effort from the IT admin, which is the motivation behind query minimization.

Our model is highly related to monotone Stochastic Boolean Function Evaluation (SBFE). There are two existing exact algorithms for SBFE that are prohibitively expensive. We propose a significantly more scalable exact algorithm. While previous exact algorithms only scale for trivial graphs (i.e., past works experimented on at most 20 edges), we empirically demonstrate that our algorithm is scalable for a wide range of much larger practical graphs (i.e., Windows domain network graphs with tens of thousands of edges).

We propose three heuristics. Our best-performing heuristic is via reducing the search horizon of the exact algorithm. The other two are via reinforcement learning (RL) and Monte Carlo tree search (MCTS). We also derive an anytime algorithm for computing the performance lower bound. Experimentally, we show that all our heuristics are near optimal. The exact algorithm-based heuristic outperforms all, surpassing RL, MCTS and 8 existing heuristics ported from SBFE and related literature.

AI-Empowered Real-Time Coordination for Wind and Energy Storage in Australian Wholesale Energy and Ancillary Service Markets

Jinhao Li

Wind energy has been increasingly adopted to mitigate climate change and support net-zero transition. However, the intermittent nature of wind generation causes wind curtailments, leading to substantial economic losses and significant waste of the excessive wind energy. To mitigate wind curtailments, the battery energy storage system (BESS) has been recognised as one of the most effective onsite backup sources, while its auxiliary role does not fully unlock the financial viability of BESS in the energy market. In particular, while reducing wind curtailments, the BESS can conduct energy arbitrage in the spot market and provide frequency control ancillary service (FCAS) in the ancillary service market. Hence, ideal BESS functioning should balance wind curtailment management and market bidding to unlock the BESS's economic value. Yet, real-time BESS operations are currently facing two major challenges: 1) coordinating wind curtailment management during their market participation; 2) mitigating uncertainties brought by the volatile energy prices and time-varying wind generation.

The Australian Energy Market Operator has acknowledged the importance of co-location of renewable energy and BESS in dedicated renewable energy zones in energy system planning, with pilot co-located wind-BESS systems built in New South Wales and South Australia. We are motivated to investigate the bidding strategy of a co-located wind-BESS system in the joint-market setting, including the spot market and the regulation FCAS market. While both facilities can simultaneously participate in the joint markets, the BESS can absorb onsite curtailed wind energy while performing energy trading, unlocking the full economic potential of the co-located system. To better characterise the coordination between the wind farm and the BESS, we decouple the wind-BESS system's market participation into two related Markov decision processes for each asset, which are optimised by the twin delayed deep deterministic policy gradient algorithm.

Using realistic wind farm data and real-time energy prices in the Australian National Electricity Market, we demonstrated the effectiveness of our coordinated bidding strategy, with bidding outcomes significantly exceeding the optimisation-based benchmark (namely predict-then-optimise) in terms of higher revenue by approximately 25%, more wind curtailment reduction by 2.3 times, and extremely faster inference computation speed. Moreover, our simulation reveals that joint-market bidding can significantly improve the financial performance of wind-BESS systems compared to participating in each market separately. Modelling results also show that using curtailed wind generation as a power source for charging the BESS can lead to additional financial gains. The successful implementation of our algorithm would encourage co-location of generation and storage assets to unlock wider system benefits.

Planning for Inertia and Resource Adequacy in a Renewable-rich Power System: A Competitive Market Expansion Approach

Amir Fayaz Heidari

We employ a competitive market expansion planning model, which is a deterministic quadratic optimization formulation, to determine the optimal investment strategies for power system technologies during a long-term planning horizon. This approach aims to address resource adequacy and security challenges within the power system, which have arisen due to the integration of a high volume of renewable energy sources. Moreover, some sensitivity analyses are carried out to quantitatively investigate the impact of uncertainty of critical input parameters on the power system planning. In our model, various electricity market players, including synchronous classical generators, renewable energy sources, and energy storage systems are categorized as either regulated or strategic. Regulated firms, categorized as price takers, aim to maximize social welfare. On the contrary, strategic players, known as price makers, prioritize their own profit using their ability to influence the market price by exercising their own market power. The expansion capacity for each of these market players is determined based on their individual long-term bidding strategies. Our model also considers several convex constraints including regional inertia requirements and operational constraints of each firm and the power system which result in a convex constraint set.

As the objective function of each firm over the convex constraint sets is convex, it can be shown that solving the collective objective functions of all individual firms is essentially tantamount to optimizing a centralized objective function. In fact, our presented centralized objective function has the same Nash Equilibrium point as the competitive game model. The centralized objective function is to maximize the total market profit which is calculated as the sum of all discounted yearly market profits over the planning horizon. Also, the market profit for each year is determined by subtracting the market inefficiency and market cost from the market income. Indeed, a strategic firm makes the market inefficient and decreases the social welfare of the market by exercising its market power. In addition, the market cost includes the investment, maintenance, and operation costs of the firms. The market income follows a quadratic function based on market demand and the market inefficiency of each strategic firm is a quadratic function of its dispatch level. On the other hand, the investment, maintenance, and operation costs for each firm are linear functions of its newly installed capacity, total capacity, and dispatch level, respectively.

Evolving Reinforcement Learning Environment to Minimize Learner's Achievable Reward: An Application on Hardening Active Directory Systems

Diksha Goel

We study a Stackelberg game between one attacker and one defender in configurable environment settings. The defender picks a specific environment configuration. The attacker observes the configuration and attacks via Reinforcement Learning (RL). The RL is trained against the observed environment. The defender's goal is to find the environment with minimum achievable reward for the attacker. We apply Evolutionary Diversity Optimization (EDO) to generate diverse population of environments for training. Environments with clearly high rewards are killed off and replaced by new offsprings to avoid wasting training time. Diversity not only improves training quality but also fits well with our RL scenario, i.e., RL agents tend to improve gradually, so a slightly worse environment earlier on may become better later. We demonstrate the effectiveness of our approach by focusing on a specific application, Active Directory (AD).

AD is the default security management system for Windows domain networks. Given its prevalence among large and small organizations worldwide, AD has become a major target by cyber attackers. An AD environment naturally describes a cyber-attack graph, where the nodes are computers/accounts/etc., and the edges represent existing accesses/known exploits that allow the attacker to gain access from one node to another. In our Stackelberg game model between one attacker and one defender on an AD graph, the attacker initially has access to a set of entry nodes. The attacker can grow this set by exploring edges. We assume that every edge has a detection rate and a failure rate. While exploring an edge, the attacker may get detected, which immediately ends the attack. The attacker may also fail to go through an edge, which does not end the attack. After encountering a failed edge, the attacker can continue the attack by exploring one of the remaining unexplored edges, originating from the same or a different source node that the attacker already has access to. There is a single destination node called the Domain Admin (DA).

The attacker aims to maximize their chance of successfully reaching DA before getting detected. The defender's task is to block a constant number of edges to decrease the attacker's chance of success. For the attacker problem, we propose a RL based policy to maximize attacker's chances of successfully reaching the DA (maximize attacker's achievable reward). We propose an Evolutionary Diversity Optimization (EDO) based defensive policy to find defensive plan configurations that minimize the attacker's success rate.

Gradient Descent Learning Algorithms for Finding Stable Fixed Points of Bimatrix Games

Kamal Mammadov

Finding Nash equilibria (or stable fixed points) in non-convex games is analogous to, but much harder than, acquiring local minima in neural nets. Even in the simplest case of two-player, two-action bimatrix games, infinitesimal gradient descent exhibits cyclic behaviour, where the action selection probabilities cycles around the Nash equilibrium without convergence. This presentation explores modifications/adjustments of gradient descent learning algorithms for finding stable fixed points of bimatrix games. We test the performance of our algorithms in the game of radar escape.

Preserving Privacy while Publishing Information

Harry McArthur

The central goal of data-dissemination agencies around the world is to collect and publish information about individuals in a population to inform important policy decisions and research objectives. Can we ensure the privacy of the individuals is maintained while retaining the utility of the data? The current privacy preserving techniques used in practice rely on injecting noise into the statistics prior to publishing. In order to quantify the level of leakage a given technique has, we can consider what an attacker can learn from the released information, or equivalently what they can reconstruct about the underlying database. Anything that can be learned about an individual is a potential violation of privacy. We discuss the design of different reconstruction methods for the purpose of informing the decision on how best to inject this noise or implement other privacy preserving techniques — optimising the utility of the data subject to certain privacy constraints. A natural property of the published statistics is to be margin-consistent, meaning the published marginals are consistent with the corresponding interior cell counts across any choice of contingency table. Generally, this property is not imposed due to tractability is- sues when simultaneously satisfying privacy constraints. We believe enforcing this property can provide significantly improved protection guarantees for the same amount of added noise, and consequently propose a novel and efficient method for generating margin-consistent noise based on Markov-Chain Monte-Carlo techniques.

Day 2, Session Three, 13:50 – 15:30

Optimizing Chance-Constrained Submodular Problems with Variable Uncertainties

Xiankun Yan

Chance constraints are frequently used to limit the probability of constraint violations in real-world optimization problems where the constraints involve stochastic components. We study chance-constrained submodular optimization problems, which capture a wide range of optimization problems with stochastic constraints. Previous studies considered submodular problems with stochastic knapsack constraints in the case where uncertainties are the same for each item that can be selected. However, uncertainty levels are usually variable with respect to the different stochastic components in real-world scenarios, and rigorous analysis for this setting is missing in the context of submodular optimization. This paper provides the first such analysis for this case, where the weights of items have the same expectation but different dispersion. We present greedy algorithms that can obtain a high-quality solution, i.e., a constant approximation ratio to the given optimal solution from the deterministic setting. In the experiments, we demonstrate that the algorithms perform effectively on several chance-constrained instances of the maximum coverage problem and the influence maximization problem.

The Chance-Constrained Travelling Thief Problem

Thilina Chathuranga Pathirage Don

The travelling thief problem (TTP) is a multi-component combinatorial optimization problem that has gained significant attention in the evolutionary computation and heuristic search literature. In our study, we introduce the chance constrained TTP which involves stochastic weights. Our problem formulation captures the stochastic aspect of the knapsack in the form of a chance constraint. Such a constraint can only be violated with a small probability. We introduce surrogate and sampling-based approaches for the chance constrained TTP in order to optimize the expected TTP score under the condition that the solution is feasible with a high probability. We use these approaches for evaluating the feasibility of solutions and incorporate our approaches into high-performing algorithms for the deterministic TTP. In our experimental investigations, we compare the performance of these algorithms and show the impact of uncertainty in connection with the underlying stochastic model.

Efficient Packing for Coles' Online Orders in Australia with Ant Colony Optimization

Bahram Farhadinia

Every day, Coles stores in Australia receive hundreds of online shopping orders, which are divided into different waves with specific timeframes for completion. To handle these orders efficiently, it's crucial to pack customer products optimally into crates. Our research team from the University of Melbourne recently partnered with an industrial project to tackle this challenge. We developed an algorithm based on Ant Colony Optimization that effectively optimizes the packing of products in this complex online order fulfillment process.

Battery on Wheel: Optimising Charging/Discharging Schedule to Reduce Emissions and Cost for EV Owners

Shizhe Zhao

Australia is the world's 13th largest emitter of overall greenhouse gas (GHG) emissions, with domestic transport accounting for approximately 17% of Australia's GHG emissions. Electric vehicles (EVs) play a vital role in achieving global climate change goals. While EVs themselves do not directly emit GHGs, they contribute to GHG emissions when charged using electricity generated from fossil fuels. Therefore, to promote urban sustainability, it is essential to encourage the use of renewable energy for charging EVs as much as possible. Furthermore, in addition to reducing system wide GHG emissions, EVs can act as large batteries, storing and then discharging surplus energy to smart grids (V2G) and premises (V2P) when needed. However, achieving this goal currently requires EV owners to make complex decisions, such as when to charge, when to discharge, and how much energy to charge or discharge while ensuring sufficient energy for their needs. Some EV owners may be reluctant to participate in this process if it does not offer them personal benefits or if it is too difficult to carry out. Hence, there is a need to make it more appealing for EV owners to use their vehicles as renewable energy storage solutions, possibly by reducing manual intervention and the operational costs of the vehicle.

In this project, we aim to automatically generate operation plans for EV owners based on their preference and ensure usability while taking into account economic (operational cost) and ecological (GHG emissions) factors. We design the project in three stages. In the early stage, we consider a simplified deterministic environment where all information is prior known, including the user's preferences, travel schedules, travel consumption, energy types, and buy/sell prices for each hour. In the next stage, we estimate realistic travel consumption based on the starting and target locations of a trip and integrate energy-aware route recommendations. In the final stage, we consider a more realistic model that the environment is non-deterministic, such as unknown user preferences, uncertain travel schedules, fluctuating renewable energy generation, real-time traffic, and dynamic energy prices.

Currently, we have completed the early stage, and the results demonstrate that in a deterministic environment, an optimal plan can reduce operational costs by 55% and GHG emissions by 73%. Additionally, it can yield a net profit of up to $55 per week when V2G is available. These results are encouraging for real-world applications and motivate us to proceed.

Benchmarking Optimisation Algorithms and Neural Network Loss Landscapes

Marcus Gallagher

Feed-forward neural networks (aka multilayer perceptrons) have been widely applied to supervised learning problems since the mid-1980s. Over this time, thousands of different datasets have been used in thousands of different experimental studies, with results reported in the literature. This research has helped to fuel tremendous progress in the field. However, documented reproduction of published experimental results has never been attempted in many cases. The availability of computational resources, software libraries and datasets creates the opportunity to attempt to reproduce and even expand on experiments that previously took a large amount of time.

In addition, it is possible to run experiments not just to try and locate single best minimizer of the training loss function, but to collect and explore numerous convergence points on a loss landscape, in order to better understand the properties of problem instances (e.g. in relation to multimodal optimization and exploratory-driven techniques such as quality-diversity search). In this talk I will discuss some results in this direction and discuss possibilities for methodology and analysis in this direction.

Day 2, Session Four, 16:00 – 17:20

Hybridising Tree-Based Pipeline Optimisation Tool with Bayesian Optimisation

Angus Kenny

The tree-based pipeline optimisation tool (TPOT) automates the building and tuning of machine learning pipelines for both classification and regression tasks. These pipelines are structured as trees, incorporating diverse data transformation and learning operations, and optimised using genetic programming methods. This talk discusses our attempts to integrate TPOT with Bayesian Optimisation (BO), aiming to bolster its capacity to explore continuous hyper-parameter domains, especially under restrictive computational constraints. We introduce and rigorously assess multiple hybrid models, such as (a) alternating or periodic employment of BO and (b) the adoption of discrete or continuous domains for BO. We gauged the efficacy of these hybrids on 6 datasets, each possessing up to 20 features and as many as 20,000 instances. While these modifications did not yield dramatically superior outcomes, they did provide some important insights into the inherent behaviours and limitations of TPOT itself, which will be instrumental in designing improved variants.

Constraint Learning Using ReLu DNN for an Optimal Bidding Engine for a Flexibility Aggregator in the European Electricity Market

Yogesh Pipada Sunil Kumar

In many real-world optimization problems, explicit relationships for constraints or objective functions may be difficult to formulate. In such cases machine learning based approaches can be used to learn such constraints and used to represent them within optimization models. However, the machine learning technique themselves might introduce non-convex constraints within the optimization model making them hard to solve.

ReLu DNNs are emerging as an interesting candidate for constraint learning approaches because they can be easily reformulated into a mixed integer linear programming problem. Additionally, these models are continuous piecewise linear and hence have the potential to represent complex non-linear relationships. However, the complexity of the MILP obtained from ReLu DNNs are strongly linked to the number of ReLu neurons. To overcome this complexity, we propose a linear reformulation of ReLu DNNs which exploit the continuous piecewise linear nature of ReLu's and can be applied to optimization problems where: 1) the output of the ReLu DNN is explicitly minimized; 2) the weights of the ReLu DNN are constrained in a specific manner. This constrained ReLu DNN trades-off accuracy for optimization problem tractability.

To demonstrate this formulation, we developed an optimal bidding engine for an aggregator of residential flexibility with a non-linear cost function in the objective. This non-linear function was represented using three approaches: piecewise linearization, ReLu DNN with unconstrained weights and ReLu DNN with constrained weights. These three formulations were implemented, and the results confirm the increased tractability of our proposed reformulation albeit with slight inaccuracy of the solution.

AI-Copilot for Business Optimisation: A Framework and A Case Study in Production Scheduling

Pivithuru Thejan

Business optimisation is an important process to help businesses gain competitive advantages by reducing operational costs, improving customer satisfaction, and mitigating risks. Although modern optimisation technologies have offered businesses different ways to formulate and solve their problems, successfully adopting these technologies still requires significant domain knowledge and optimisation expertise. In this research, we consider business optimisation from a computational and mathematical perspective where one tries to minimize or maximize an important characteristic of a process by an appropriate choice of decisions.

Recently, Large Language Models (LLMs) have become increasingly popular due to their broad applications. Initiated by transformer for machine translation, LLMs have been quickly adopted within different software and business functions. In fact, the motivation behind this research is to leverage code-generating LLMs to automate problem formulation. Therefore, the time and effort needed in problem formulation can be minimized. However, automating problem formulations is a non trivial task for code-generating LLMs due to complex constraints, different optimisation requirements, and the need of selecting the most suitable optimisation technique. Additionally, due to token limitation, code-generating LLMs cannot generate large problem formulations. Also, the large computational and memory requirements of some code-generating LLMs limit practical use. In addition, the existing performance evaluation metrics in code-generating LLMs are not appropriate for problem formulation since the result as well as the optimisation technique need to be considered.

Therefore, we introduce our AI-Copilot as a step towards automating problem formulations for complex real-world optimisation problems. To do so, we select production scheduling as a case study as it has been comprehensively researched in the past and contains complex constraints and different optimisation objectives. We fine-tune a code-generating LLM, which uses limited memory, and computational resources, using a data set created by us that comprised 100 pairs of problem descriptions and formulations. As a result, we minimize the requirement of large training data, rather than training an LLM from scratch for problem formulation. In addition, we apply modularization and prompt engineering techniques on our AI-Copilot to cater to the token limitation when formulating complex problem formulations. Furthermore, we use loss and execution-based performance evaluation metrics to assess the accuracy and quality of problem formulations compared to existing evaluation metrics. The experiment results demonstrate that we can synthesize complex and large problem formulations for a typical business optimisation problem in production scheduling.

Learning to Branch in Integer Programming

Edward Lam

This talk gives a quick overview of a reinforcement learning algorithm for learning to branch in integer programming. The task is to select a variable currently taking fractional values in the linear relaxation of an integer program such that a small branch-and-bound tree is produced. The algorithm came in third place at the Machine Learning for Combinatorial Optimization (ML4CO) competition at NeurIPS 2021.