Sessions are subject to change and will be finalised together

	Monday	Tuesday	Wednesday	Thursday (till 1 pm)
9:00 – 10:15	Welcome Words (30 mins)+Intro talks (2 hours)	Brandon Amos + Sofia Michel + Ernestine Grossmann	Fredrik Manne + Ali Vakilian + Michele Lombardi	Sandra Kiefer + Marianne Defresne + Nysret Musliu
10:15-10:45	Coffee Break	Coffee Break	Coffee Break	Coffee Break
10:45 – 12:00	Remaining Intro talks	Quentin Cappart + Dolores Romero Morales + + Misha Khodak	Jens Schloeter + Alexandre Forel + Louis-Martin Rousseau	Maurice Wenig + Open Problem Session 3 + Final words
12:15 – 13:00	Lunch	Lunch	Lunch	Lunch
13:30-15:10	Tias Guns + Adam Polak	Open Problem Session 1	Open Problem Session 2	
15:15 – 16:00	Coffee and Cake	Coffee and Cake	Coffee and Cake	
16:00 – 17:40	Bistra Dilkina + Deepak Ajwani	Summary from Open Problem Session + + Alexander Lindermayer + Aaron Ferber + Sylvie Thiebaux	Summary from Open Problem Session + Hike!!	

Note: Talks are 25 minute duration each. This can be in the form of 20 minutes talk + 5 minutes questions or 15 minute talk + 10 minute question.

Tutorials

Presenter: Tias Guns
Title: Decision-focused Learning
Abstract: Increasingly, combinatorial optimisation problems follow a predict + optimize paradigm, where part of the parameters (costs, volumes, capacities) are predicted from data, and those predictions are fed into a downstream combinatorial optimisation problem. How to best train these predictive models? Decision-focused learning (DFL) is an emerging paradigm in machine learning which trains a model to optimize decisions, integrating prediction and optimization in an end-to-end system. This paradigm holds the promise to revolutionize decision-making in many real-world applications which operate under uncertainty, where the estimation of unknown parameters within these decision models often becomes a substantial roadblock to high-quality solutions. This talk presents a

comprehensive review of DFL. It provides an in-depth analysis of the various techniques devised to integrate machine learning and optimization models, introduces a taxonomy of DFL methods distinguished by their unique characteristics, and conducts an extensive empirical evaluation of these methods proposing suitable benchmark dataset and tasks for DFL. Finally, we'll provide valuable insights into current and potential future avenues in DFL research.

Presenter: Adam Polak
Title: Approximation Algorithms with Predictions
Abstract: I will talk about utilizing predictions to improve over approximation guarantees of classic algorithms, without increasing the running time. I will cover prior results going in that direction, and then I will talk about our recent work in which we propose a generic method for a wide class of optimization problems that ask to select a feasible subset of input items of minimal (or maximal) total weight. This gives simple (near-)linear-time algorithms for, e.g., Vertex Cover, Steiner Tree, Min-Weight Perfect Matching, Knapsack, and Clique. Our algorithms produce an optimal solution when provided with perfect predictions and their approximation ratio smoothly degrades with increasing prediction error. With small enough prediction error, we achieve approximation guarantees that are beyond the reach without predictions in given time bounds, as exemplified by the NP-hardness and APX-hardness of many of the above problems. Although we show our approach to be optimal for this class of problems as a whole, there is a potential for exploiting specific structural properties of individual problems to obtain improved bounds; we demonstrate this on the Steiner Tree problem. I will conclude with an empirical evaluation of our approach. This is based on joint work with Antonios Antoniadis, Marek Eliáš, and Moritz Venzin.

Presenter: Bistra Dilkina
Title: Learning for Combinatorial Optimisation
Abstract: TBA

Presenter: Deepak Ajwani
Title: Learning problems in Algorithm design and engineering
Abstract: In this talk, we will explore the learning problems arising in algorithm design, engineering and analysis. We will discuss (i) ways in which the insights from the algorithm engineering literature can be used in problem-specific learning solutions, (ii) how learning techniques can be applied to fill the underspecified parts in algorithms (e.g., ordering in Bellman-Ford algorithm), (iii) how machine learning can help in identifying and estimating the various parameters used in the design of parameterised algorithms, (iv) how learning techniques can be used for the design of interpretable heuristics and (v) if learning techniques can be used for finding combinatorial structures and worst-case inputs for algorithm analysis.

Talks

Presenter: Brandon Amos
Title: On LLM prompt optimization and amortization
Abstract: Prompting LLMs is a challenging art where different ways of expressing the same idea can lead to drastically different responses. Not prompting in the right way gives suboptimal performance and has led to the budding space of prompt engineering and optimization techniques. A standard example here is that appending the string "let's think step by step" to the prompt may significantly improve the quality on few-shot classification tasks. In this session, we'll first cover how prompt optimization can be expressed as a combinatorial optimization problem (over the token sequence space) and overview the standard methods here. (A warning for this audience: standard combinatorial solvers or approaches are often not used, and instead approaches are specialized to LLMs.) Beyond this, prompt optimization problems are often not solved a single time in isolation, but are repeatedly solved for every new task or piece of information we want to extract from the language model. So, we'll conclude with an overview of learned optimization, or amortization, to share the information between the repeatedly-solved optimization problems.

Presenter: Sofia Michel
Title: Neural Combinatorial Optimization – One Model to Rule them All?
Abstract: I presented recent works in NCO with a focus on generalization. I started with an overview of standard constructive NCO approaches. Then, I introduced the BQ-MDP framework and its strong out-of-distribution generalization. Next, I presented the idea of a generalist combinatorial optimization agent learning (GOAL), showing how a unique pretrained model can be efficiently fine-tuned to solve a variety of unseen CO problems.

Presenter: Ernestine Grossmann
Title: Screening for Data Reductions in Combinatorial Optimization: A Graph Neural Network Approach
Abstract: Combinatorial optimization problems play a pivotal role in various domains, challenging researchers to find optimal solutions within large solution spaces. One key strategy to tackle the complexity of these problems is data reduction, where instances are simplified to facilitate more efficient algorithms. Even though there are numerous reduction rules available, they are often not used in applications since the time needed to test them for the whole graph exhaustively becomes infeasible. With this, a new problem in the form of early reduction screening emerges. In this talk, we tackle this problem in the context of maximum weight independent sets using graph neural networks. Before checking if an expensive data reduction rule is applicable, we consult a GNN oracle to decide if the probability of successfully reducing the graph is sufficiently large. With this approach, we can use even expensive reduction rules successfully in feasible time. We introduce a new supervised learning dataset and provide first results using established graph neural network architectures.

Presenter: Quentin Cappart
Title: Enhancing constraint programming with machine learning: current challenges and future opportunities
Abstract: Constraint programming (CP) is a well-established method for tackling combinatorial problems. Traditionally, CP has focused on solving isolated problem instances, often overlooking the fact that these instances frequently originate from related data distributions. In recent years, there has been a growing interest in leveraging machine

learning, particularly neural networks, to enhance CP solvers by utilizing historical data. Despite this interest, it remains unclear how to effectively integrate learning into CP engines to boost overall performance. In this presentation, I will share my experience in tackling this challenge, from my initial attempts to my current research directions.

Presenter: Dolores Romero Morales
Title: Group Counterfactual Explanations with a Mathematical Optimization Lens
Abstract: Counterfactual Analysis has shown to be a powerful tool in the burgeoning field of Explainable Artificial Intelligence. In Supervised Classification, this means associating with each record a so-called counterfactual explanation: an instance that is close to the record and whose probability of being classified in the opposite class by a given classifier is high. In this talk we take a stakeholder perspective, and we address the setting in which a group of counterfactual explanations is sought for a group of instances. We introduce some mathematical optimization models as illustration of each possible allocation rule between counterfactuals and instances and tasks beyond Supervised Classification.

Presenter: Misha Khodak
Title: Efficiently learning instance-optimal linear system solvers
Abstract: Augmenting classical algorithms with learned predictions or configurations has found many successful applications in energy management, database systems, scientific computing, and beyond. At the same time it has been theoretically challenging to go beyond the computationally inefficient learning of static predictors and configurations. We study the problem of sequentially solving linear systems, a fundamental problem in numerical computing, and introduce an algorithm that efficiently learns to do instance-optimal linear solver configuration while using only the number of iterations as feedback. Our approach points towards new directions for analyzing iterative methods, designing surrogate objectives for optimizing algorithmic costs, and incorporating tools from online learning into algorithm design.

Presenter: Marianne Defresne
Title: Learning to solve real-world puzzles: from Sudoku to protein design
Abstract: Real-life decision making often involves reasoning on ill-defined problems, where exact constraints or parameters (such as costs) are unknown. The goal of Decision-Focused Learning (DFL) is to automatize the definition of problem parameters by using Deep Learning to extract knowledge out of the environment. The main challenge here is to combine discrete optimization for reasoning with continuous optimization for learning.
I will present one method to learn discrete graphical models, the reasoning framework we chose. We first assessed it on learning the rules of Sudoku, a popular benchmark for DFL methods. We then apply it to Computational Protein Design, a real-world problem that can be described and tackled with graphical models. We deep learned the interactions within existing proteins to better guide the design towards new proteins of interest.

Presenter: Alexander Linermayr
Title: Non-Clairvoyant Scheduling with Predictions
Abstract: In this talk, we explore some recent and ongoing advancements in non-clairvoyant scheduling in the learning-augmented framework, which integrates error-prone predictions into online algorithm design. We examine various prediction models, showcasing algorithms for non-clairvoyant scheduling, also in a more general setting with online precedence constraints, with strong error-dependent performance guarantees. In particular, we ask which type of information is required and how an algorithm should receive it to achieve reasonable performance guarantees. Moreover, we consider recently popular prediction models for

other problems, and discuss how they may be integrated into scheduling problems.

Presenter: Sylvie Thiebaux
Title: Graph Learning for Planning
Abstract: I will present recent work on graph representation learning to guide the search of AI planners. I will introduce GNN and other graph learning representations that exploit the relational structure of planning domains. They allow our planner GOOSE to learn search guidance (e.g. heuristic cost estimates, state rankings) from solutions to just a few small problems, and solve substantially larger problems than trained on. Perhaps surprisingly, our experimental results show that classical machine learning approaches vastly outperform deep learning ones in this context. Moreover, Greedy Best-First Search guided by our best learnt heuristics outperforms the state of the art model-based planner, Lama, on the problems of the latest International Planning Competition Learning track, leading to the possibility that learnt heuristics may replace existing model-based heuristics in the near future.

Presenter: Fredrik Manne
Title: Graph Neural Networks as Ordering Heuristics for Parallel Graph Coloring
Abstract: The graph coloring problem asks for an assignment of the minimum number of distinct colors to vertices in an undirected graph with the constraint that no pair of adjacent vertices share the same color. The problem is a thoroughly studied NP-hard combinatorial problem with several real-world applications. As such, a number of greedy heuristics have been suggested that strike a good balance between coloring quality, execution time, and also parallel scalability. In this work, we introduce a graph neural network (GNN) based ordering heuristic and demonstrate that it outperforms existing greedy ordering heuristics both on quality and performance. Previous results have demonstrated that GNNs can produce high-quality colorings but at the expense of excessive running time. The current paper is the first that brings the execution time down to compete with existing greedy heuristics. Our GNN model is trained using both supervised and unsupervised techniques. The experimental results show that a 2-layer GNN model can achieve execution times between the largest degree first (LF) and smallest degree last (SL) ordering heuristics while outperforming both on coloring quality. Increasing the number of layers improves the coloring quality further, and it is only at four layers that SL becomes faster than the GNN. Finally, our GNN- based coloring heuristic achieves superior scaling in the parallel setting compared to both SL and LF.

Presenter: Ali Vakilian
Title: Learning-Based Algorithms for Graph Searching Problems
Abstract. We consider the problem of graph searching with prediction recently introduced by [Banerjee et al., ITCS'23]. In this problem, an agent, starting at some vertex r has to traverse a (potentially unknown) graph G to find a hidden goal node g while minimizing the total distance travelled. We study a setting in which at any node v, the agent receives a noisy estimate of the distance from v to g. We design algorithms for this search task on unknown graphs. We establish the first formal guarantees on unknown weighted graphs and provide lower bounds showing that the algorithms we propose have optimal or nearly-optimal dependence on the prediction error. Further, we perform numerical experiments demonstrating that in addition to being robust to adversarial error, our algorithms perform well in typical instances in which the error is stochastic. Finally, we provide alternative simpler performance bounds on the algorithms of [Banerjee et al., ITCS'23] for the case of searching on a known graph, and establish new lower bounds for this setting

Presenter: Michele Lombardi
Title: Progress and Open Questions on DFL and Constrained NNs
Abstract: This talk will provide perspectives, progress, and open questions on two distinct forms of integration between optimization and learning. First, we will present an analysis of Decision Focused Learning approaches, highlighting structural problem properties that can cause their advantage w.r.t. classical techniques to be significantly diminished. We will suggest a possible way out, with interesting practical applications, together with one solution technique.
Second, we consider the problem of enforcing hard constraints in Neural Networks, discussing a training-time approach with satisfaction guarantees. The method combines Projected Gradient Descent with a neural architecture that embeds a trainable over-approximation module. For both the considered setting, we highlight open problems and promising research directions.

Presenter: Jens Schloeter
Title: Two-Oracle Models for Matroid Optimization Problems
Abstract: In many combinatorial optimization problems, access to input data is modeled in an abstract way via oracles, e.g., independence oracles for matroids, distance functions in a metric space, or comparators for sorting elements. Depending on the underlying application, the actual use of such oracles can be computationally expensive, as it may require access to slow database systems or large machine learning (ML) models. In the two-oracle model, we assume access to a second, weaker oracle, which is computationally cheap but may provide inaccurate answers. The goal is to design algorithms that exploit the information provided by the cheap oracle to minimize the number of calls to the expensive oracle.
Two-oracle models have been studied in several ML applications, such as data labeling with weak and strong data labelers, text clustering with two oracles for text similarity, or more generally in ML pipelines that combine the use of scalable and expensive models. More recently, Bai and Coester [NeurIPS'23] considered two-oracle models from the point of view of learning-augmented algorithm design for sorting with access to two comparators. As usual, the goal is to design robust algorithms with performance guarantees that improve with the quality of the second, cheap oracle. In this talk, we discuss the two-oracle model and its applications, and study it from the perspective of learning-augmented algorithms by applying it to the fundamental problem of finding a maximum-weight basis in a matroid.
The talk is based on a joint work with Franziska Eberle, Felix Hommelsheim, Alexander Lindermayr, Nicole Megow, and Zhenwei Liu that has been accepted for poster presentation at NeurIPS'24.

Presenter: Alexandre Forel
Title: The differentiable feasibility pump
Abstract: The feasibility pump algorithm is a widely used heuristic to find feasible primal solutions to mixed-integer linear problems. Many extensions of the algorithm have been proposed. Yet, its core algorithm remains centered around two key steps: solving the linear relaxation of the original problem to obtain a solution that respects the constraints, and rounding it to obtain an integer solution. This paper shows that the feasibility pump and many of its follow-ups can be seen as gradient-descent algorithms with specific parameters. A central aspect of this reinterpretation is observing that the traditional algorithm differentiates the solution of the linear relaxation with respect to its cost. This reinterpretation opens many opportunities for improving the performance of the original algorithm. We study how to modify the gradient-update step as well as extending its loss function. We perform extensive experiments on instances from the MIPLIB library and show that these modifications can substantially reduce the number of iterations needed to find a primal solution.

Presenter: Louis-Martin Rousseau
Title: Optimize then predict: an imitation-based learning framework.
Abstract: Predictive models are playing an increasingly pivotal role in optimizing decision-making. This talk introduces an approach wherein a series of (possibly) stochastic sequential decision problems are initially solved offline, and subsequently, a predictive model is constructed based on the offline solutions to facilitate real-time decision-making. The applicability of this methodology is demonstrated through two use cases: patient radiation therapy, where managers need to preserve capacity for emergency cases, and a novel dynamic employee call-timing problem for the scheduling of casual personnel for on-call work shifts.

Presenter: Maurice Wenig
Title: Serving MPE Queries on Tensor Networks by Computing Derivatives
Abstract: In probabilistic graphical models, lots of inference queries can be formulated as computations on tensor networks. Computing the partition function of the graphical model boils down to summing over a product of the factors of the graphical model, which is a very common tensor network computation. When we maximize over the product of the factors instead of summing, which can also be done in tensor networks, then a maximum probability is computed instead of the normalizing constant. MPE Queries now ask for not only the maximum probability, but also an assignment of variables that achieves this probability.
In this presentation we explore how to compute such an assignment using derivatives of this maximizing tensor network computation.

Presenter: Sandra Kiefer
Title: The Expressive Power of Graph Neural Networks
Abstract: Graph Neural Networks (GNNs) are a machine learning architecture to learn functions on graphs. For example, since problem instances for combinatorial optimisation tasks are often modelled as graphs, GNNs have recently received attention as a natural framework for finding good heuristics in neural optimisation approaches.
The question which functions can actually be learnt by message-passing GNNs and which ones exceed their power has been studied extensively. In this talk, I will consider it from a graph-theoretical perspective. I will survey the Weisfeiler-Leman algorithm as a combinatorial procedure to analyse and compare graph structure, and I will discuss some results concerning the power of the algorithm on natural graph classes. The findings directly translate into insights about the power of GNNs and of their extensions to higher-dimensional neural networks.

Presenter: Nysret Musliu
Title: Machine Learning for Algorithm Selection
Abstract: Algorithm selection addresses the challenge of determining which of the available algorithms is most appropriate for solving a particular problem instance. Hyper-heuristics are a high-level, problem-independent approach that aims to automate the design of heuristic methods by combining and selecting low-level heuristics.

In this talk, we will present our work on algorithm selection for various problem domains, including scheduling and tree decomposition. We will also discuss our recent work on reinforcement learning for selection hyper-heuristics.