LION 17

June 4 - 8, 2023 Nice, France

Book of abstracts

(Version 17.05.23)

Program chairs: Meinolf Sellmann (InsideOpt) and Kevin Tierney (Bielefeld University)

333 Hyperparameter Optimization and Weight Initializations: A Performance Analysis for Deep Learning Models

Nicki Lena Kämpf

With the advent of black-box Deep Learning Models (DLM), the importance of hyperparameter optimization has grown rapidly. This is because the performance of these DLM heavily depends on the appropriate hyperparameters. However, one aspect has been neglected when it comes to the hyperparameter set: the choice of the weight initializations. Different studies by Gloriot and Bengio or He et al. have shown, that the convergence and thus the performance of DLMs is related to the weight initializations. Research studies in this field have resulted in a variety of initialization schemes. To give some perspective, there are currently 14 different weight initialization schemes implemented in the pytorch package, and 13 different weight initialization schemes are implemented in the keras package. The weight initialization schemes can again be adjusted by adapting the random states that influence which initial weight values are drawn from the distributions of the underlying weight initialization scheme. The goal of this study is to research the effects of weight initialization schemes and random states for different DLM architectures. Due to the numerous combinations of weight initialization schemes and random states, a Sequential Model Based Optimization (SMBO) algorithm will be used to tune the hyperparameters, the weight initializations and the random states jointly. The motivation for SMBO is that it has proven to mitigate the trade-off between the run-time of the hyperparameter optimization and the performance of the models. The first promising results in time-series forecasting with Temporal Fusion Transformers revealed that the DLM's performance varies greatly when different weight initializations and random states are used. Only for specific weight initialization schemes and random states the DLM outperforms simpler models. In conclusion, the integration of weight initializations and random states for hyperparameter optimization algorithms is advocated for.

568 Anomaly Classification to Enable Self-Healing in Cyber Physical Systems using Process Mining

Uphar Singh, Deepak Gajjala, Rahamatullah Khondoker, Harshit Gupta, Ayush Sinha and O.P. Vyas

Industrial Cyber Physical Systems (CPS) are large-scale critical infrastructures that are vulnerable to cyberattacks with wide-ranging consequences. Being a combination of heterogeneous devices and protocols, the large scale CPS anomaly is also exposed to critical vulnerabilities. These vulnerabilities are treated in terms of anomalies and cyberattacks, and their detection and corresponding self-healing mechanisms on large-scale critical infrastructures can be challenging because of their massive size and interconnections. With the objective of process optimization through anomaly detection and conformance checking approach, the present work addresses different issues, such as event log generation with tools such as PLG 2.0 for data-driven approach. The self-healing is enabled through anomaly classification ensemble learning-based machine learning models. The work uses process mining for analyzing event log files, and then combinations of conformance-checking methods with ensemble classification models were used to best classify anomalies. Finally, the proposed work establishes that, in comparison to techniques like KNN and C-SVC, the proposed ensemble models perform better, with an accuracy of 84.7% using trace alignment as a conformance technique with gradient boosting to classify anomalies, with the end objectives of process improvement.

613 Hyper-box classification model using mathematical programming

Georgios I. Liapis and Lazaros G. Papageorgiou

Classification constitutes focal topic of study within the machine learning research community. Interpretable machine learning algorithms have been gaining ground against black box models because people want to understand the decision-making process. Mathematical programming based classifiers have received attention because they can compete with state-of-the-art algorithms in terms of accuracy and interpretability. This work introduces a single-level hyper-box classification approach, which is formulated mathematically as Mixed Integer Linear Programming model. Its objective is to identify the patterns of the dataset using a hyper-box representation. Hyper-boxes enclose as many samples of the corresponding class as possible. At the same time, they are not allowed to overlap with hyper-boxes of different class. The interpretability of the approach stems from the fact that IF-THEN rules can easily be generated. Towards the evaluation of the performance of the proposed method, its prediction accuracy is compared to other state-of-the-art interpretable approaches in a number of real-world datasets. The results provide evidence that the algorithm can compare favourably against well-known counterparts.

728 A leak localization algorithm in water distribution networks using probabilistic leak representation and optimal transport distance

Andrea Ponti, Ilaria Giordani, Antonio Candelieri and Francesco Archetti

Leaks in water distribution networks are estimated to account for up to 30% of the total distributed water: the increasing demand, and the skyrocketing energy cost have made leak localization and adoption ever more important to water utilities. Each leak scenario is run on a simulation model to compute the resulting values of pressure and flows over the whole network. The values recorded by the sensors are seen as the features of one leak scenario and can be considered as the signature of the leak. The key distinguishing element in the present paper is the representation of a leak signature, over the simulation horizon, as a discrete probability distribution. In this representation the similarity between leaks can be captured by a distance between their associated probability distributions. This maps the problem of leak detection from the Euclidean physical space into a space whose elements are histograms: this space is structured by a distance between histograms, namely the Wasserstein distance also known as the optimal transport distance. This choice also matches the physics of the system: indeed, the equations modelling the generation of flow and pressure data are non-linear. Non-linear data structure is better represented by the Wasserstein distance than by the Euclidean distance. The signatures obtained through the simulation of a large set of leak scenarios are non-linearly clustered according in the Wasserstein space using Wasserstein barycenters as centroids. As a new set of sensor measurements arrives, the related signature is associated to the cluster with the closest barycenter. The location of the simulated leaks belonging to that cluster are the possible locations of the observed leak. This new theoretical and computational framework allows a richer representation of pressure and flow data embedding both the modelling and the computational modules in a space, the Wasserstein space, whose elements are the histograms, endowed with the Wasserstein distance. The computational experiments on benchmark and real-world networks confirm the feasibility of the proposed approach.

864 Fast and Robust Constrained Optimization via Evolutionary and Quadratic Programming

Konstantinos Chatzilygeroudis and Michael Vrahatis

Many efficient and effective approaches have been proposed in the evolutionary computation literature for solving constrained optimization problems. Most of the approaches assume that both the objective function and the constraints are black-box functions, while a few of them can take advantage of the gradient information. On the other hand, when the gradient information is available, the most versatile approaches are arguably the ones coming from the numerical optimization literature. Perhaps the most popular methods in this field are sequential quadratic programming and interior point. Despite their success, those methods require accurate gradients and usually require a well-shaped initialization to work as expected. In the paper at hand, a novel hybrid method, named UPSO-QP, is presented that is based on particle swarm optimization and borrows ideas from the numerical optimization literature and sequential quadratic programming approaches. The proposed method is evaluated on numerous constrained optimization tasks from simple low dimensional problems to high dimensional realistic trajectory optimization scenarios, and showcase that is able to outperform other evolutionary algorithms both in terms of convergence speed as well as performance, while also being robust to noisy gradients and bad initialization.

1124 Bayesian Optimization for Function Compositions with Applications to Dynamic Pricing

Kunal Jain, Prabuchandran K.J. and Tejas Bodas

Bayesian Optimization (BO) is used to find the global optima of black box functions. In this work, we propose a practical BO method of function compositions where the form of the composition is known but the constituent functions are expensive to evaluate. By assuming an independent Gaussian process (GP) model for each of the constituent black-box function, we propose Expected Improvement (EI) and Upper Confidence Bound (UCB) based BO algorithms and demonstrate their ability to outperform not just vanilla BO but also the current state-of-art algorithms. We demonstrate a novel application of the proposed methods to dynamic pricing in revenue management when the underlying demand function is expensive to evaluate.

1131 Sensorimotor Learning with Stability Guarantees via Autonomous Neural Dynamic Policies

Dionis Totsila, Konstantinos Chatzilygeroudis, Dimitrios Kanoulas and Ioannis Hatzilygeroudis

State-of-the-art sensorimotor learning algorithms, either in the context of reinforcement learning or imitation learning, offer policies that can often produce unstable behaviors, damaging the robot and/or the environment. Moreover, it is very difficult to interpret the optimized controller and analyze its behavior and/or performance. Traditional robot learning, on the contrary, relies on dynamical system-based policies that can be analyzed for stability/safety. Such policies, however, are neither flexible nor generic and usually work only with proprioceptive sensor states. In this work, we bridge the gap between generic neural network policies and dynamical system-based policies, and we introduce Autonomous Neural Dynamic Policies (ANDPs) that: (a) are based on autonomous dynamical systems, (b) always produce asymptotically stable behaviors, and (c) are more flexible than traditional stable dynamical system-based policies. ANDPs are fully differentiable, flexible generic-policies that can be used for both imitation learning and reinforcement learning setups, while ensuring asymptotic stability. Through several experiments, we explore the flexibility and capacity of ANDPs in several imitation and reinforcement learning tasks including experiments with image observations. The results show that ANDPs combine the benefits of both neural network-based and dynamical system-based methods.

1173 A Bayesian optimization algorithm for constrained simulation optimization problems with heteroscedastic noise

Sasan Amini and Inneke Van Nieuwenhuyse

In this research, we develop a Bayesian optimization algorithm to solve expensive, constrained problems. We consider the presence of heteroscedastic noise in the evaluations and thus propose a new acquisition function to account for this noise in the search for the optimal point. We use stochastic kriging to fit the metamodels, and we provide computational results to highlight the importance of accounting for the heteroscedastic noise in the search for the optimal solution. Finally, we propose some promising directions for further research.

1175 Hierarchical Machine Unlearning

Zhu Hongbin, Xia Yuxiao, Li Yunzhao, Li Wei, Liu Kang and Gao Xianzhou

In recent years, deep neural networks have enjoyed tremendous success in industry and academia, especially for their applications in visual recognition and natural language processing. While large-scale deep models bring incredible performance, their massive data requirements pose a huge threat to data privacy protection. With the growing emphasis on data security, the study of data privacy leakage in machine learning, such as machine unlearning, has become increasingly important. There have been many works on machine unlearning, and other research has proposed training several sub-models to speed up the retraining process, by dividing the training data into several disjoint fragments. When the impact of a particular data point in the model is to be removed, the model owner simply retrains the sub-model containing this data point. Nevertheless, current learning methods for machine unlearning are still not widely used due to model applicability, usage overhead, etc. Based on this situation, we propose a novel hierarchical learning method, Hierarical Machine Unlearning (HMU), with the known distribution of unlearning requests. Compared with previous methods, ours has better efficiency. Using the known distribution, the data can be partitioned and sorted, thus reducing the overhead in the data deletion process. We propose to train the model using the hierarchical data set after partitioning, which further reduces the loss of prediction accuracy of the existing methods. It is also combined with the methods of incremental learning to speed up the training process. Finally, the effectiveness and efficiency of the method proposed in this paper are verified by multiple experiments.

1371 The Dynamic RORO Stowage Planning Problem

Alastair Main, Filipe Rodrigues and Dario Pacino

Over the past few years, significant attention has been given to reducing the shipping industry's greenhouse gas emissions (GHG). One area of research that has received particular focus is stowage planning for Roll-on/Roll-off (RORO) vessels. RORO shipping is dedicated to the transportation

of wheeled cargo, e.g., cars, semi-trailers, buses, and farm equipment. A known way for ships to reduce GHG emissions is by increasing sailing time and utilizing slow steaming.
Efficient stowage plans will reduce vessel turnaround time and increase the sailing time for vessels in ports. The handling of cargo follows a FILO queue which can result in blocking cargo when loading and unloading the vessel at the port. When assigning a placement of the cargo in the vessel's deck, it is essential to consider paths taken by the cargo on board the vessel. Furthermore, all current research assumes that all the cargo is available for stowage. This is rarely the case, as cargo arrives at the port over time while stowing is in operation. Therefore, the stowage planning process must also consider cargo arrival times.
We propose a novel mathematical model that utilizes a weighted objective function to address these challenges. This function seeks to minimize the relationship between fuel consumption cost from sailing time and revenue gained from shipping cargo. By scheduling the cargo loading sequence, we aim to reduce the time spent handling and re-handling cargo at each port. Our study focuses on a single deck layout for a vessel calling multiple ports, and we will present the results of our mathematical model and accompanying metaheuristic.

1669 Explaining the Behavior of Reinforcement Learning Agents using Association Rules

Zahra Parham, Vi Tching de Lille and Quentin Cappart

Deep reinforcement learning algorithms are increasingly used to drive decision-making systems. However, there exists a known tension between the efficiency of a machine learning algorithm and its level of explainability. Generally speaking, increased efficiency comes with the cost of decisions that are harder to explain. This concern is related to explainable artificial intelligence, which is a hot topic in the research community. In this paper, we propose to explain the behaviour of a black-box sequential decision process, built with a deep reinforcement learning algorithm, thanks to standard data mining tools, i.e. association rules. We apply this idea to the design of playing bots, which is ubiquitous in the video game industry. To do so, we designed three agents trained with a deep Q-learning algorithm for the game Street FighterTurbo II. Each agent has a specific playing style and the data mining algorithm aims to find rules maximizing the lift, while ensuring a minimum threshold for the confidence and the support. Experiments show that association rules can provide insights on the behavior of each agent, and reflect their specific playing style. We believe that this work is a next step towards the explanation of complex models in deep reinforcement learning.

1766 Deep Randomized Networks for Fast Learning

Richárd Rádli and László Czúni

Deep learning neural networks show a significant improvement over shallow ones in complex problems. Their main disadvantage is their memory requirements, the vanishing gradient problem, and the time consuming solutions to find the best achievable weights and other parameters. Since many applications (such as continuous learning) would need fast training, one possible solution is the application of sub-networks which can be trained very fast. Randomized single layer networks became very popular due to their fast optimization while their extensions, for more complex structures, could increase their prediction accuracy. In our paper we show a new approach to build deep neural models for classification tasks with an iterative, pseudo-inverse optimization technique. We compare the performance with a state-of-the-art backpropagation method and the best known randomized approach called hierarchical extreme learning machine. Computation time and prediction accuracy are evaluated on 12 benchmark datasets, showing that our approach is competitive in many cases.

2111 Real-world process discovery from low-level event data

Stephane Cholet, Franck Lefebure and Cecile Thuault

New perspectives on monitoring and optimising business processes have emerged from advances in process mining research, as well as techniques for their discovery, conformance checking and enhancement. Many identify process discovery as the most challenging task, as it is the first in line, and as the methods deployed could impact the effectiveness of subsequent analysis. More and more companies are ready to use process mining products, but not at any cost. They require them to be highly reliable, fast and to introduce a limited overload on their resources (including human). In a real-word setting, specific business constraints add to common issues, such as high data frequency and asynchrony. This also adds to the complexity of real-world processes and of process mining itself. Mainstream studies propose to use a finite set of high-level and business-oriented event Hyperparameter Optimization and Weight Initializations: A Performance Analysis for Deep Learning Modelslogs, where key attributes (i.e., case, activity and timestamp) are known, and apply unscaled discovery techniques to produce control-flow process models. In this research, we propose an original approach we have designed and deployed to mine processes of businesses. It features fully streamed and real-time techniques to mine low-level technical event data, where key attributes do not exist and have to be forged. We will focus on the scope of process discovery, and expose our adoption of an organizational perspective, driven by (semi-) unsupervised discovery, streaming and scaling features.

2087 Generative models via Optimal Transport and Gaussian Processes

Antonio Candelieri, Andrea Ponti and Francesco Archetti

Generative models have recently gained a renewed interest due to their success in the development of new real-life applications, such as artificial intelligence generated images, texts, audios. The most recent and successful approaches combine neural network learning and Optimal Transport theory, exploiting the so-called transportation map/plan to generate a new element of a domain starting from an element of a different one, while preserving statistical properties of the data generation processes of the two domains. Although effective, the Neural Optimal Transport (NOT) approach is largely computationally expensive – due to the training of two nested deep neural networks – and requires injecting additional noise to improve generative properties. In this paper we present an alternative method, based on Gaussian Process (GP) regression, which overcomes these limitations. Contrary to a neural model, a GP is probabilistic, meaning that, for a given input, it provides both a prediction and the associated uncertainty. Thus, the generative properties are, by design, guaranteed by sampling the generated element around the prediction and depending on the uncertainty. Results on both toy examples and a dataset of images are provided to empirically demonstrate the benefits of the proposed approach.

2088 Multi-Agent Reinforcement Learning for Strategic Bidding in Two Stage Electricity Markets

Francesco Morri, Hélène Le Cadre, Luce Brotcorne and Pierre Gruet

Our goal is to study the dynamics of electricity markets involving multiple competitive generators through multi-agent reinforcement learning (MARL) approaches. We start by formulating the electricity market as a two-stage stochastic game, involving a finite set of conventional and renewable energy producers, which bid on the day-ahead market, and an Independent System Operator (ISO), which is responsible for the clearing of the market. We assume that a predetermined part of the producers are non-strategic, bidding at their marginal costs, while the others might bid strategically trying to learn the outcome of the clearing. In the first stage, the strategic producers optimize simultaneously their bids by minimizing their expected costs (opposite of the expected profits), which is the difference between their pro- duction cost and the payment they receive from the ISO. The renewable energy producers' objective functions include a penalty assigning a cost to the imbalances caused by their forecast errors. In the second stage, the ISO receives the bids of all the producers. It clears the market by determining the activated volumes for each producer and a price minimizing the total cost under capacity constraints, including a conditional value at risk (CVaR) constraint for the renewable producers, capturing the risk aversion level that the requested volume violates their uncertain capacity. We derive closed form expressions for the producers' best-responses considering pay-as-clear and pay-as-bid as pricing schemes, and simulate the market dynamics, using MARL. To that purpose, we rely on modified versions of two actor-critic algorithms: Deep Deterministic Policy Gradient and Soft Actor-Critic. The simulations show how the producers adapt dynamically their strategies to learn the best bidding strategy, under limited information exchange. Finally, we identify conditions for the convergence of MARL algorithms to local equilibria of the stochastic game.

2957 GPU for Monte Carlo Search

Lilian Buzer and Tristan Cazenave

Monte Carlo Search algorithms give excellent results for some combinatorial optimization problems and for some games. They can be parallelized efficiently on high-end CPU servers. Nested Monte Carlo Search is an algorithm that parallelizes well. We take advantage of this property to obtain large speedups running it on low cost GPUs. The combinatorial optimization problem we use for the experiments is the Snake-in-the-Box. It is a graph theory problem for which Nested Monte Carlo Search previously improved lower bounds. It has applications in electrical engineering, coding theory, and computer network topologies. Using a low cost GPU, we obtain speedups as high as 420 compared to a single CPU.

2963 Learning the Bias Weights for Generalized Nested Rollout Policy Adaptation

Julien Sentuc, Farah Ellouze, Jean-Yves Lucas and Tristan Cazenave

Generalized Nested Rollout Policy Adaptation (GNRPA) is a Monte Carlo search algorithm for single player games and optimization problems. In this paper we propose to modify GNRPA in order to automatically learn the bias weights. The goal is both to obtain better results on sets of dissimilar instances, and also to avoid some hyperparameters settings. Experiments show that it improves the algorithm for two different optimization problems: the Vehicle Routing Problem and 3D Bin Packing.

3301 Heuristics selection with ML in CP Optimizer

Hugues Juille, Renaud Dumeur and Paul Shaw

IBM® ILOG® CP Optimizer (CPO) is a constraints solver that integrates multiple heuristics with the goal of handling a large diversity of combinatorial and scheduling problems while exposing a simple interface to users. CPO's intent is to enable users to focus on problem modelling while automating the configuration of its optimization engine to solve the problem. For that purpose, CPO proposes an Auto search mode which implements a hard-coded logic to configure its search engine based on the runtime environment and some metrics computed on the input problem. This logic is the outcome of a mix of carefully designed rules and fine-tuning using experimental benchmarks.

This paper explores the use of Machine Learning (ML) for the off-line configuration of CPO solver based on an analysis of problem instances. This data-driven effort has been motivated by the availability of a proprietary database of diverse benchmark problems that is used to evaluate and document CPO performance before each release. This work also addresses two methodological challenges: the ability of the trained predictive models to robustly generalize to the diverse set of problems that may be encountered in practice, and the integration of this new ML stage in the development workflow of the CPO product. Overall, this work resulted in a speedup improvement of about 14% (resp. 31%) on Combinatorial problems and about 5% (resp. 6%) on Scheduling problems when solving with 4 workers (resp. 8 workers), compared to the regular CPO solver.

3464 Model-based feature selection for neural networks: A mixed-integer programming approach

Shudian Zhao, Calvin Tsay and Jan Kronqvist

In this work, we develop a novel input feature selection framework for ReLU-based deep neural networks (DNNs), which builds upon a mixed-integer optimization approach. While the method is generally applicable to various classification tasks, we focus on finding input features for image classification for clarity of presentation. The idea is to use a trained DNN, or an ensemble of trained DNNs, to identify the salient input features. The input feature selection is formulated as a sequence of mixed-integer linear programming (MILP) problems that find sets of sparse inputs that maximize the classification confidence of each category. These "inverse" problems are regularized by the number of inputs selected for each category and by distribution constraints. Numerical results on the well-known MNIST and FashionMNIST datasets show that the proposed input feature selection allows us to drastically reduce the size of the input to $\sim 15\%$ while maintaining a good classification accuracy. This allows us to design DNNs with significantly fewer connections, reducing computational effort and producing DNNs that are more robust towards adversarial attacks.

3602 An Error-Based Measure for Concept Drift Detection and Characterization

Antoine Bugnicourt, Riad Mokadem, Franck Morvan and Nadia Bebeshina

Continual learning is an increasingly studied field, aiming at regulating catastrophic forgetting for online machine learning tasks. In this article, we propose a prediction error measure for continual learning, to detect concept drift induced from learned data input before the learning step. In addition, we check this measure's ability for characterization of the drift. For these purposes, we propose an algorithm to compute the proposed measure on a data stream while also estimating concept drift. Then, we calculate the correlation coefficients between this estimate and our measurement, using time series analysis. To validate our proposal, we base our experiments on simulated streams of metadata collected from an industrial dataset corresponding to real conversation data. The results show that the proposed measure constitutes a reliable criterion for concept drift detection. They also show that a characterization of the drift relative to components of the stream is possible thanks to the proposed measure.

3715 Predict, Tune and Optimize for Data-Driven Shift Scheduling with Uncertain Demands

Michael Römer, Felix Hagemann and Till Porrmann

When it comes to data-driven optimization under uncertainty, it is well-known that a classical predict-then-optimize pipeline in which point forecasts maximizing predictive accuracy are plugged into a deterministic optimization model typically leads to a poor expected performance. In stochastic programming, one aims at obtaining better expected performance by explicitly representing the joint probability distribution in the optimization model, e.g. in form of a sample approximation. A downside of that approach is that it gives rise to large-scale model instances that are hard to solve, in particular in a combinatorial optimization setting. An alternative approach that recently attracted considerable interest is to rely on "decision-focused" point predictions to be fed into a deterministic optimization model. This approach, referred to as decision-focused learning or predict and optimize in the literature, aims to train prediction models in a way that the expected performance of the decisions obtained with the (prediction-informed) deterministic model is maximized. In this paper, we propose to generalize this idea by optimizing not only parameters affecting the prediction but also additional parameters influencing other (non-stochastic) parts of the optimization model. We propose to simultaneously optimize both types of parameters with the goal of maximizing expected decision quality and refer to this approach as predict, tune and optimize. We demonstrate the usefulness of the approach for a multi-activity shift scheduling problem under uncertainty in which almost 1000 demand parameters are affected by uncertainty. Specifically, we show that while decision-oriented tuning of point forecasts usually yields better results than a simple predict-then-optimize approach, adding the possibility to modify additional parameters considerably improves the expected performance which becomes competitive with a stochastic programming approach.

3727 Machine Learning for Combinatorial Optimisation of Partially-Specified Problems: Regret Minimisation as a Unifying Lens

Stefano Teso, Laurens Bliek, Andrea Borghesi, Michele Lombardi, Neil Yorke-Smith, Tias Guns and Andrea Passerini

It is increasingly common to solve combinatorial optimisation problems that are partially specified. We survey the case where the objective function or the relations between variables are not known or are only partially specified, while the hard constraints are assumed to be given. Such problems are commonplace in, for example, vehicle routing (driver preferences), multi-objective optimisation (weight penalties), and energy-aware scheduling (energy prices). One challenge is to learn the unspecified parts from available data, while taking into account the set of hard constraints that a solution must satisfy. Furthermore, solving the optimisation problem (especially during learning) is often computationally very demanding. This work overviews four seemingly unrelated approaches, that can each be viewed as learning the objective function of a hard combinatorial optimisation problem: 1) decision-focused learning ('predict + optimise'), 2) structured-output prediction, 3) empirical model learning and 4) surrogate-based optimisation. We formalise each learning paradigm, at first in the ways commonly found in the literature, and then bring the formalisations together in a compatible way using the notion of regret. This leads to a unified combinatorial optimisation problem that is a generalisation of the different optimisation problems used in the four discussed approaches. A regret minimisation approach is formulated for solving the unified problem with the four discussed approaches, each with their own assumptions but all in one compatible formulation. We also highlight the opportunities for cross-fertilisation, such as the use of bi-level formulations in structured-output prediction, or using iterative acquisition strategies in empirical model learning. Finally, we survey open research directions such as scalability issues and incorporating prior knowledge.

3745 Automated packing systems

David Alvarez Martinez, Daniel Giraldo, Ana Maria Montes Franco, Daniel Cuellar-Usaquen, German Fernando Pantoja Benavides, Ruben Iván Bolaños, Carlos Rodriguez, Valentina Bedoya and Juan Martínez

E-commerce has disrupted traditional business models and the supply chains that support them. Distributors have had to solve even more complex problems of warehousing, storage, and distribution of goods due to the sheer volume of packages they must handle. Distributors have turned to robot-supported or fully automated operations, such as loading and unloading goods, where machines achieve outstanding performance. This paper reviews automated three-dimensional packing problems and identifies the main features considered in these applications. In the review performed on the solution methods of the packing problem, the gaps to exploit in future research are shown. In this work, a learning-based optimization algorithm used in a real-world robotic packing cell is proposed and compared with classical packing heuristics. The results obtained with a set of instances based on historical information from a well-known regional distribution company show that the proposed algorithm outperforms classical heuristics. In future work, it is proposed to consider constraints associated with the kinematics of the manipulators.

3921 On Learning When to Decompose Graphical Models

Aleksandra Petrova and Javier Larrosa

Decomposition is a well-known algorithmic technique for Graphical Models. It is commonly believed that such a technique is cost-effective for instances with low width. In this paper, we show on a large data set of real-life inspired instances that this is not the case. To better understand this result, we narrow our study and consider k-tree instances where the width is well controlled and get similar results. Finally, we show that adding a few simple features and using simple Machine Learning models we can predict the convenience to decompose with an accuracy of more than 85%, which produces time reductions in standard benchmarks of nearly 90%.

4002 Repositioning Fleet Vehicles: a Learning Pipeline

Augustin Parjadis, Quentin Cappart, Quentin Massoteau and Louis-Martin Rousseau

Managing a fleet of vehicles under uncertainty requires careful planning and adaptability. We consider a ride-hailing problem where the operator manages vehicle repositioning to maximize responsiveness. This paper introduces a supervised learning pipeline that uses past trip data to reposition vehicles while adapting to fleet activity, a geographical zone, and seasonal or daily request variation. The pipeline incorporates trip features, such as medical motives of transportation for ambulances and the time and location of the trips. This provides a better estimate of the probability that a given vehicle will be required in a particular sector and provides insights into which events and trip features should be incorporated into the decision-making process for better fleet management and improved reactivity. This tool has been developed for, and used by, operators of an ambulance company in Belgium. Using predictors for ambulance repositioning reduces at least 10reaction distance.

4133 Bayesian Decision Trees Inspired from Evolutionary Algorithms

Efthyvoulos Drousiotis, Alexander Phillips, Paul Spirakis and Simon Maskell

Bayesian Decision Trees (DTs) are generally considered a more advanced and accurate model than a regular Decision Tree (DT) because they can handle complex and uncertain data. Existing work on Bayesian DTs uses Markov Chain Monte Carlo (MCMC) with an accept-reject mechanism and sample using naive proposals to proceed to the next iteration, which can be slow because of the burn-in time needed. We can reduce the burn-in period by proposing a more sophisticated way of sampling or by designing a different numerical Bayesian approach. In this paper, we propose a replacement of the MCMC with an inherently parallel algorithm, the Sequential Monte Carlo

(SMC), and a more effective sampling strategy inspired by the Evolutionary Algorithms (EA). Experiments show that SMC combined with the EA can produce more accurate results compared to MCMC in 100 times fewer iterations.

4238 Towards Tackling MaxSAT by Combining Nested Monte Carlo with Local Search

Hui Wang, Abdallah Saffidine and Tristan Cazenave

Recent work proposed the UCTMAXSAT algorithm to address Maximum Satisfiability Problems (MaxSAT) and shown improved performance over pure Stochastic Local Search algorithms (SLS). UCTMAXSAT is based on Monte Carlo Tree Search but it uses SLS instead of purely random playouts. In this work, we introduce two algorithmic variations over UCTMAXSAT. We carry an empirical analysis on MaxSAT benchmarks from recent competitions and establish that both ideas lead to performance improvements. First, a nesting of the tree search inspired by the Nested Monte Carlo Search algorithm is effective on most instance types in the benchmark. Second, we observe that using a static flip limit in SLS, the ideal budget depends heavily on the instance size and we propose to set it dynamically. We show that it is a robust way to achieve comparable performance on a variety of instances without requiring additional tuning.

4488 Optimizing Product Offerings with Deep Reinforcement Learning

Chaher Alzaman

The ability to manage product inventory effectively is a critical factor for the success of sellers of short-life products. In today's highly competitive market, it is essential for these businesses to have a streamlined and efficient approach to selecting the products they stock. This paper proposes a novel solution to address the challenge faced by these sellers. Specifically, we focus on the problem of selecting a bundle of products, or product line, to stock in cost-competitive environment. To accomplish this, we use deep reinforcement learning to train agents to make decisions based on historical data patterns where agents continuously improve their product selection policy over time. Grocery stores are a prime example of organizations that must carefully select their product offerings in real-time. With numerous choices available, they must choose which products to keep in stock based on factors such as customer demand, seasonality, and, more importantly, shelf-life. Second, small to medium-sized sellers often rely on retail giants' e-platforms to sell specific products (dropshipping). For instance, traders may study data from Amazon to gain insights into which products are outperforming the market average and adjust their product offerings accordingly. This selection process can be quite complex as traders must consider factors such as customer preferences, market trends, and sales data.

We use deep reinforcement learning methodology that combines deep neural networks and reinforcement learning algorithms to train agents to make decisions based on historical data. Specifically, we train agents using Markovian chains to identify the best-performing products and adapt our decision-making strategies in response to changing market conditions and customer demand. This approach enables the agents to continuously learn and improve the product selection policies over time, taking into account the complex and dynamic environment of the marketplace.

4555 Relational Graph Attention-based Deep Reinforcement Learning: An Application to Flexible Job Shop Scheduling with Sequence-dependent Setup Times

Amirreza Farahani, Martijn Van Elzakker, Laura Genga, Pavel Troubil and Remco Dijkman

This paper tackles a manufacturing scheduling problem using an Edge Guided Relational Graph

Attention-based Deep Reinforcement Learning approach. Unlike state-of-the-art approaches, the proposed method can deal with machine flexibility and sequence dependency of the setup times in the Job Shop Scheduling Problem. Furthermore, the proposed approach is size-agnostic. We evaluated our method against standard priority dispatching rules based on data that reflect a realistic scenario, designed on the basis of a practical case study at the Dassault Systèmes company. We used an industry-strength large neighborhood search based algorithm as benchmark. The results show that the proposed method outperforms the priority dispatching rules in terms of makespan, obtaining an average makespan difference with the best tested priority dispatching rules of 4.45% and 12.52%.

4466 Expert-Iteration for Combinatorial Optimization

Arnaud Sors, Darko Drakulic, Florian May, Sofia Michel and Jean-Marc Andreoli

Deep Learning has been successfully applied to learn heuristics for Combinatorial Optimization (CO) problems [Bengio et al, 2021]. While most Reinforcement Learning based approaches for CO use model-free algorithms [Mazyavkina et al, 2020], we propose to explore the application of Generalized Policy Iteration methods [Sutton and Barto 2018] as model-based approaches for deterministic CO problems. Examples of these methods include AlphaZero [Silver et al, 2017] and MuZero [Schrittwieser et al, 2019], that give state-of-the-art results for games (Go, Chess, Atari, etc). Other variants have been successfully applied to power grid congestion management [Dorfer et al, 2022], video compression [Mandhane et al, 2022] and robot control [Schrittwieser et al, 2021]. In this presentation, we focus on Expert-Iteration (ExIt) [Anthony et al, 2017] which was proposed concurrently with AlphaZero for a broader class of problems. We share our experience and current work on applying ExIt to solving combinatorial optimization problems. We illustrate our results on the Travelling Salesman Problem (TSP) as one of the simplest combinatorial optimization problems, yet representative of a larger class of problems. First, we present our neural model, which is a Transformer model [Vaswani et al 2017] that accounts for the symmetries of the TSP. Then, we present our training procedure with ExIt and provide experimental results. Finally, we share our experience on the challenges and solutions for implementing ExHyper-box classification model using mathematical programmingIt in practice; esp. parallel tree search, multi-machine training, ease of implementation versus speed tradeoff — with the idea of answering the question: how to best implement ExIt for your CO problem?

4570 Managing Perishable Inventory with Distributional Forecasts

Mark Velednitsky and Philip Cerles

This work presents a new continuous-review policy for managing the inventory of produce in grocery stores. There are well-established inventory management policies for single order periods (ex. newsvendor) or items that can be held indefinitely without spoilage (ex. EOQ). We consider the intermediate case: the shelf lives of our products are typically 3-14 days and orders are typically placed every 2-4 days. Thus, items last for several periods but not indefinitely.
Our new policy, ShrinkAware, optimizes an objective function that includes terms inspired by the classic newsvendor model and the classic EOQ model. We compare the performance of our policy to an incumbent policy, MI-Policy (minimum inventory policy).
ShrinkAware consumes a demand forecast, generated by a deep learning model, as well as shelf life estimates. When evaluating the performance of ShrinkAware versus MI-Policy, we explain why ShrinkAware is outperforming MI-Policy and the implications for upstream ML systems.
Three of our key findings are:

(1) Using a full distributional forecast is significantly more powerful than a point forecast, even a point forecast calculated at a few predetermined quantiles of the distribution. (2) Using a safety stock proportional to the forecast rather than the available shelf space improves the policy's performance considerably. (3) The shelf life is very important when it is less than three re-ordering cycles. After that, the holding cost term is dominant.
We implemented ShrinkAware in over 250 grocery stores in a major US chain, resulting in a 6% decrease in inventory lost to shrink, a 4% decrease in out-of-stock lost sales, and a 2% increase in overall profits compared to MI-Policy. Previous experiments have shown that MI-Policy reduces shrink relative to human decision-makers by up to 25%.

4836 Experimental Digital Twin for Job Shops with Transportation Agents

Aymen Gannouni, Luis Felipe Casas Murillo, Marco Kemmerling, Anas Abdelrazeq and Robert H. Schmitt

Production scheduling in multi-stage manufacturing environments is subject to combinatorial optimization problems, such as the Job Shop Problem (JSP). The transportation of materials when assigned to mobile agents, such as Automated Guided Vehicles (AGVs), results in a Job Shop Problem with Transportation Agents (JSPTA). The transportation tasks require routing the AGVs within the physical space of the production environment. Efficient scheduling of production and material flow is thus crucial to enable flexible manufacturing systems. Neural combinatorial optimization has evolved to solve combinatorial optimization problems using deep Reinforcement Learning (RL). The key aim is to learn robust heuristics that tackle the trade-off of optimality versus time complexity and scale better to dynamic changes in the problem. The present simulation environments used to train RL agents for solving the JSPTA lack accessibility (e.g. use of proprietary software), configurability (e.g. changing shop floor layout), and extendability (e.g. implementing other RL methods). This research aims to address this gap by designing an Experimental Digital Twin (EDT) for the JSPTA. It represents an RL environment that considers the physical space for the execution of production jobs with transportation agents. We created our EDT using a simulation tool selected based on requirement analysis and tested it with a customized state-of-the-art neural combinatorial approach against two common Priority Dispatching Rules (PDRs). With a focus on the makespan, our findings reveal that the neural combinatorial approach outperformed the other PDRs, even when tested on unseen shop floor layouts. Furthermore, our results call for further investigation of multi-agent collaboration and layout optimization. Our EDT is a first step towards creating self-adaptive manufacturing systems and testing potential optimization scenarios before transferring them to real-world applications.

4968 Learning with Neural Networks to solve Resource Leveling Problems

Lena Wohlert and Jürgen Zimmermann

Neural networks and (deep) reinforcement learning are increasingly used for heuristics to solve combinatorial optimization problems. Corresponding research in project scheduling has focused on the resource-constrained project scheduling problem, where a schedule generation scheme is extended to learn priority values of activities depending on instance features and previous scheduling steps. In our work, we address the related resource leveling problem, which involves not only deciding on the scheduling sequence of the activities but also their start times to balance the resource utilization over the planning horizon. For this problem, we implemented an actor-critic algorithm to determine scheduling priorities for the activities, followed by a greedy start time selection. The learned policy is used to sample an activity, using the negative increase in the objective function value when scheduling an activity as a reward. We are also working on how to learn a

meaningful probability distribution over the possible start times of each activity depending on the current partial schedule. In the future, we want to incorporate research on other combinatorial optimization problems. Promising topics are the use of recurrent neural networks or the integration of a (similar) deep reinforcement learning framework in e.g. local search or genetic algorithms. We are confident that our research will support the successful use of neural networks for many other project scheduling problems.

5071 Early Detection of Emerging Phishing Threats: a Process Mining Approach

Krutika Vyas, Nachiket Tapas and O. P. Vyas

According to IBM's X-Force Threat Intelligence Index 2023 report, phishing remains the most common infection vector, accounting for 41% of incidents in 2022. Additionally, IBM's cost of data breach report in 2021 revealed that phishing was the second costliest cyber attack on organizations, resulting in losses of $4.65 million. Although AI/ML approaches have been employed to detect phishing attacks, their effectiveness is limited by the specific model deployed. These approaches do not address zero-day exploitation through phishing. To address this issue, the proposed work explores the use of process mining to detect phishing attacks within organizations. By collecting event log data on user behavior online, process mining can detect anomalies via user's actions. This method allows for the identification of the actual sequence of events that occur during a phishing attack and can prevent zero-day exploitation.

5088 Learning to Prune Electric Vehicle Routing Problems

James Fitzpatrick, Deepak Ajwani and Paula Carroll

Electric vehicle variants of vehicle routing problems are significantly more difficult and time-consuming to solve than traditional variants. Many solution techniques fall short of the performance that has been achieved for traditional problem variants. Machine learning approaches have been proposed as a general end-to-end heuristic solution technique for routing problems. These techniques have so far proven flexible but don't compete with traditional approaches on well-studied problem variants. However, developing traditional techniques to solve electric vehicle routing problems is time-consuming. In this work we extend the learning-to-prune framework to the case where exact solution techniques cannot be used to gather labeled training data. We propose a highly-adaptable deep learning heuristic to create high-quality solutions in reasonable computational time. We demonstrate the approach to solve electric vehicle routing with nonlinear charging functions. We incorporate the machine learning heuristics as elements of an exact branch-and-bound matheuristic, and evaluate performance on a benchmark dataset. The results of computational experiments demonstrate the usefulness of our approach from the point of view of variable sparsification.

5339 A matheuristic approach for electric bus fDeep Randomized Networks for Fast Learningleet scheduling

Raka Jovanovic, Sertac Bayhan and Stefan Voss

In recent years, there has been an increasing growth in the number of electric vehicles on the road. An important part of this process is the electrification of public transport with the use of electric buses. There are several differences between scheduling an electric or diesel bus fleet to cover a public transport timetable. The main reason for this is that electric buses have a shorter range and need to be charged during operating hours. The related optimization problems are often modeled using mixed-integer programming (MIP). An issue is that standard MIP solvers usually cannot solve problem instances corresponding to real-world applications of the model within a reasonable time limit. In this paper, this is addressed by extending the fixed set search to a matheuristic setting. The conducted computational experiments show that the new approach can be applied to much larger problems than the basic MIP. In addition, the proposed approach significantly outperforms other heuristic and metaheuristic methods on the problem of interest for problem instances up to a specific size.

5489 Surrogate Membership for Inferred Metrics in Fairness Evaluation

Melinda Thielbar, Serdar Kadioglu, Chenhui Zhang, Rick Pack and Lukas Dannull

As artificial intelligence becomes more embedded into daily activities, it is imperative to ensure models perform well for all sub-groups. This is particularly important when models include underprivileged populations. Binary fairness metrics, which compare model performance for protected groups to the rest of the model population, are an important way to guard against unwanted bias. However, a significant drawback of these binary fairness metrics is that they require protected group membership attributes. In many practical scenarios, protected status for individuals is sparse, unavailable, or even illegal to collect. This paper extends binary fairness metrics from deterministic membership attributes to their surrogate counterpart under the probabilistic setting. We show that it is possible to conduct binary fairness evaluation when exact protected attributes are not immediately available but their surrogate as likelihoods is accessible. Our inferred metrics calculated from surrogates are proved to be valid under standard statistical assumptions. Moreover, we do not require the surrogate variable to be strongly related to protected class membership; inferred metrics remain valid even when membership in the protected and unprotected groups is equally likely for many levels of the surrogate variable. Finally, we demonstrate the effectiveness of our approach using publicly available data from the Home Mortgage Disclosure Act and simulated benchmarks that mimic real-world conditions under different levels of model disparity.

5711 Enabling Flexible Grid Management through Intelligent Demand Response Scheduling and Neural Network-based NILM in Residential Microgrids

Mohamed Saâd El Harrab and Michel Nakhla

The increase of renewable energy sources presents new challenges for power grid operators due to their unpredictable, highly variable and uncontrollable nature. Smart Microgrids provide a flexible solution by enabling the use of demand response (DR) mechanisms, allowing efficient power grid management through adaptive energy consumption adjustments. We suggest an optimal demand response scheduling approach for a residential Microgrid. For each household, we use a Deep Learning-based Non-Intrusive Load Monitoring (NILM) system, incorporating LSTM, DAE, and SGN techniques, to identify transferable loads and forecast their consumption. This allows us to assess their micro-flexibilities potential. Then, an optimal scheduling strategy is determined by maximizing the microgrid operator's DR payoff while ensuring load demand. Simulation results indicate that the proposed DR scheduling contributes to a reduction in the overall operating costs while promoting a more efficient utilization of energy resources.

5732 The BeMi Stardust: a Structured Ensemble of Binarized Neural Networks

Ambrogio Maria Bernardelli, Simone Milanesi, Stefano Gualandi and Hoong Chuin Lau

Binarized Neural Networks (BNNs) are receiving increasing attention due to their lightweight architecture and ability to run on low-power devices, given the fact that they can be implemented using Boolean operations. The state-of-the-art for training classification BNNs restricted to few-shot learning is based on a Mixed Integer Programming (MIP) approach. This paper proposes the BeMi ensemble, a structured architecture of classification-designed BNNs based on training a single BNN for each possible pair of classes and applying a majority voting scheme to predict the final output. The training of a single BNN discriminating between two classes is achieved by a MIP model that optimizes a lexicographic multi-objective function according to robustness and simplicity principles. This approach results in training networks whose output is not affected by small perturbations on the input and whose number of active weights is as small as possible, while good accuracy is preserved. We computationally validate our model using the MNIST and Fashion-MNIST datasets using up to 40 training images per class. Our structured ensemble outperforms both BNNs trained by stochastic gradient descent and state-of-the-art MIP-based approaches. While the previous approaches achieve an average accuracy of 51.1% on the MNIST dataset, the BeMi ensemble achieves an average accuracy of 61.7% when trained with 10 images per class and 76.4% when trained with 40 images per class.

5963 Discovering Explicit Scale-Up Criteria in Crisis Response with Decision Mining

Britt Lukassen, Laura Genga and Yingqian Zhang

In modern society, incidents such as road incidents or fires, occur daily, requiring institutions to develop appropriate management protocols to react quickly. When an incident requires coordination among different emergency services or it is estimated to have a significant impact on the population, crisis management processes are used. The overall aim of crisis management is to provide the right resources to manage the incident and return to a normal situation as soon as possible. However, the decision to scale up an incident to a crisis level is often left to the experience of operational commanders without explicit criteria or guidelines. In this research, we propose a framework combining data-driven decision-mining approaches with implicit knowledge formalization techniques to discover explicit criteria to support decision-makers in crisis response. We tested our approach in a case study at VRU, the safety region for the region Utrecht, in The Netherlands. The obtained results show that the approach has been able to extract criteria acknowledged by the decision-makers, which is the first step to developing appropriate guidelines to steer the decisional process of the incident scaling up.

7001 Job Shop Scheduling via Deep Reinforcement Learning: a Sequence to Sequence approach

Giovanni Bonetta, Davide Zago, Rossella Cancelliere and Andrea Grosso

Job scheduling is a well-known Combinatorial Optimization problem with endless applications. Well planned schedules bring many benefits in the context of automated systems: among others, they limit production costs and waste. Nevertheless, the NP-hardness of this problem makes it essential to use heuristics whose design is difficult, requires specialized knowledge and often produces methods tailored to the specific task. This paper presents an original end-to-end Deep Reinforcement Learning approach to scheduling that automatically learns dispatching rules. Our technique is inspired by natural language encoder-decoder models for sequence processing and has never been used, to the best of our knowledge, for scheduling purposes. We applied and tested our method in particular to some benchmark instances of Job Shop Problem, but this technique is general enough to be potentially used to tackle other different optimal job scheduling tasks with minimal intervention. Results demonstrate that we outperform many classical approaches exploiting priority dispatching rules and show competitive results on state-of-the-art Deep Reinforcement Learning ones.

7091 CLS-Luigi: A Framework for Automated Synthesis and Execution of Decision Pipelines

Anne Meyer, Jan Bessai, Hadi Kutabi and Daniel Scholtyssek

To automate production and logistics decision-making, different data processing steps–ranging from simple data pre-processing over machine learning to optimization algorithms–are combined. These decision automation systems can be implemented by decision pipelines, which interconnect algorithmic processing steps in directed acyclic graphs. However, there is not a single best algorithm for each step in the pipeline, and, hence, there is not a single best pipeline. We introduce CLS-Luigi, a framework capable of combining the above-mentioned processing steps to solve the pipeline creation problem for decision pipelines.
CLS-Luigi combines a type-theoretic framework for software component synthesis (Combinatory Logic Synthesizer, CLS) with Luigi, a tool to build and execute pipelines for batch jobs. CLS manages the consistent creation of (all) pipeline variants based on a repository of nodes (components), guaranteeing soundness and completeness. Luigi schedules the execution of pipelines and optimizes resource usage with caching to avoid rerunning identical sub-pipelines. CLS-Luigi allows natural modeling by specifying the necessary input to a node based on the types of predecessor nodes. Nodes in the repository (e.g., type classifier or optimizer) are reusable for new pipelines. The repository is easily expandable for Python programmers by implementing problem-specific (or new problem-agnostic) nodes. Hyper-parameter tuning is covered by providing sets of promising configurations for each node. To our knowledge, there are no projects of similar scope.
In this presentation, we introduce the current state of CLS-Luigi and its basic ideas. We elaborate on the need for the automated creation of decision pipelines as a new research field and present our framework as a first step toward progress in it. Special emphasis is given to the advantages of a holistic treatment of entire pipelines that overcomes limitations of prior work directed at single parts.

7213 Generating a Graph Colouring Heuristic with Deep Q-Learning and Graph Neural Networks

George Watkins, Giovanni Montana and Juergen Branke

The graph colouring problem consists of assigning labels, or colours, to the vertices of a graph such that no two adjacent vertices share the same colour. In this work we investigate whether deep reinforcement learning can be used to discover a competitive construction heuristic for graph colouring. Our proposed approach, ReLCol, uses deep Q-learning together with a graph neural network for feature extraction, and employs a novel way of parameterising the graph that results in improved performance. Using standard benchmark graphs with varied topologies, we empirically evaluate the benefits and limitations of the heuristic learned by ReLCol relative to existing construction algorithms, and demonstrate that reinforcement learning is a promising direction for further research on the graph colouring problem.

7633 Multi-Task Predict-then-Optimize

Bo Tang and Elias Khalil

The predict-then-optimize framework arises in a wide variety of applications where the unknown cost coefficients of an optimization problem are first predicted based on contextual features and then used to solve the problem. In this work, we extend the predict-then-optimize framework to a multi-task setting: contextual features must be used to predict cost coefficients of multiple optimization problems, possibly with different feasible regions, simultaneously. For instance, in a vehicle dispatch/routing application, features such as time-of-day, traffic, and weather must be used to predict travel times on the edges of a road network for multiple traveling salesperson problems that span different target locations and multiple s - t shortest path problems with different source-target pairs. We propose a set of methods for this setting, with the most sophisticated one drawing on advances in multi-task deep learning that enable information sharing between tasks for improved learning, particularly in the small-data regime. Our experiments demonstrate that multi-task predict-then-optimize methods provide good tradeoffs in performance among different tasks, particularly with less training data and more tasks.

8137 Improving Subtour Elimination Constraint Generation in Branch-and-Cut Algorithms for the TSP with Machine Learning

Thi Quynh Trang Vo, Viet Hung Nguyen, Paul Weng and Mourad Baiou

Branch-and-Cut is a widely-used method for solving integer programming problems exactly. In recent years, researchers have been exploring ways to use Machine Learning to improve the decision-making process of Branch-and-Cut algorithms. While much of this research focuses on selecting nodes, variables, and cuts [11,9,26], less attention has been paid to designing efficient cut generation strategies in Branch-and-Cut algorithms, despite its large impact on the algorithm performance. In this paper, we focus on improving the generation of subtour elimination constraints, a core and compulsory class of cuts in Branch-and-Cut algorithms devoted to solving the Traveling Salesman Problem, which is one of the most studied combinatorial optimization problems. Our approach takes advantage of Machine Learning to address two questions before launching the separation routine to find cuts at a node of the search tree: 1) Do violated subtour elimination constraints exist? 2) If yes, is it worth generating them? We consider the former as a binary classification problem and adopt a Graph Neural Network as a classifier. By formulating subtour elimination constraint generation as a Markov decision problem, the latter can be handled through an agent trained by reinforcement learning. Our method can leverage the underlying graph structure of fractional solutions in the search tree to enhance its decision-making. Furthermore, once trained, the proposed Machine Learning model can be applied to any graph of any size (in terms of the number of vertices and edges). Numerical results show that our approach can significantly accelerate the performance of subtour elimination constraints in Branch-and-Cut algorithms for the Traveling Salesman Problem.

8347 Learn, Compare, Search: One Sawmill's Search for the Best Cutting Patterns Across And/or Trees

Marc-André Ménard, Michael Morin, Mohammed Khachan, Jonathan Gaudreault and Claude-Guy Quimper

A sawmilling process scans a wood log and must establish a series of cutting and rotating operations to perform in order to obtain the set of lumbers having the most value. The search space can be expressed as an and/or tree. Providing an optimal solution, however, may take too much time. The complete search for all possibilities can take several minutes per log and there is no guarantee that a high-value cut for a log will be encountered early in the process. Furthermore, sawmills usually have several hundred logs to process and the available computing time is limited. We propose to learn the best branching decisions from previous wood logs and define a metric to compare two wood logs in order to branch first on the options that worked well for similar logs. This approach (Learn, Compare, Search, or LCS) can be injected into the search process, whether we use a basic Depth-First Search (DFS) or the state-of-the-art Monte Carlo Tree Search (MCTS). Experiments were carried on by modifying an industrial wood cutting simulator. When computation time is limited to five seconds, LCS reduced the lost value by 47.42% when using DFS and by 17.86% when using MCTS.

8862 Reinforcement Learning for routing in mixed-shelves warehouses

Laurin Luttmann and Lin Xie

In recent years, machine learning (ML) algorithms for combinatorial optimization (CO) problems received a surge of attention. While early approaches failed to outperform traditional OR methods, the gap between handcrafted and learned heuristics has been steadily closiHeuristics selection with ML in CP Optimizerng. Especially the employment of more advanced network architectures and learning algorithms as well as their combination with search strategies like beam search have led to on-par results for many CO problems like TSP and VRP. Yet, the performance aspect is not the only reason why OR and ML researchers deem neural CO (NCO) promising. One of the main advantages of using ML to solve CO problems is the reduction in running time compared to heuristics. Many applications require quick solutions obtained within reasonable time frames which cannot be provided by extensive search algorithms and exact solvers. One such application is picker routing in warehouses, where short-term changes due to sick leave or canceled customer requests might require fast changes to the current plan. Modern warehouses are often designed in a mixed-shelves fashion, where items of the same stock keeping unit (SKU) are spread over the shelves, leading to multiple storage positions per SKU. Pickers can then decide between these alternative positions when collecting the items demanded by customers. Current NCO literature does not cover any VRP variant that incorporates the selection of alternative storage positions. Therefore, in this work we develop a transformer-based neural network which we train using the REINFORCE algorithm to obtain short picker routes in mixed-shelves warehouses. For small test instances for which we were still able to generate exact solutions, our method finds optimal or near-optimal solutions within seconds. We therefore contribute to the field of NCO by applying this emerging technology to a new problem which is more complex than the standard problems used for benchmarking in existing NCO literature.

8888 Algorithm Configuration in the UPF: Exploiting capabilities of Selector

Dimitri Weiß, Elias Schede, Kevin Tierney

The Unified Planning Framework (UPF) offers a convenient way to work with automated planning technology by providing an abstraction layer. It allows users to define planning problems in a planning engine agnostic manner, giving them the flexibility to utilize the planning engines available on their system. However, selecting the optimal parameters for a planning engine poses a significant challenge, as these engines have numerous adjustable parameters that directly affect their performance. Manually adjusting parameters of an algorithm requires extensive domain knowledge and a substantial time investment. Algorithm Configuration (AC) techniques have become commonplace in addressing this issue. These techniques aim to optimize the performance of parameterized algorithms, with objectives ranging from reducing the time required to find a solution to enhancing the quality of the generated solutions. AC methods have shown effectiveness in various problem domains, including automated planning, by finding high-quality parameter configurations. In this study, we propose the use of the ensemble-based AC approach selector for automatically configuring planning engines. We evaluate the impact of the AC approach on the performance of the planning engine lpg using problem instances included in the UPF.

9167 Dynamic Police Patrol Scheduling with Multi-Agent Reinforcement Learning

Songhan Wong, Waldy Joe and Hoong Chuin Lau

Effective police patrol scheduling is essential in projecting police presence and ensuring readiness in responding to unexpected events in urban environments. However, scheduling patrols can be a challenging task as it requires balancing between two conflicting objectives namely projecting presence (proactive patrol) and incident response (reactive patrol). This task is made even more challenging with the fact that patrol schedules do not remain static as occurrences of dynamic incidents can disrupt the existing schedules. In this paper, we propose a solution to this problem using Multi-Agent Reinforcement Learning (MARL) to address the Dynamic Bi-objective Police Patrol Dispatching and Rescheduling Problem (DPRP). Our solution utilizes an Asynchronous Proximal Policy Optimization-based (APPO) actor-critic method that learns a policy to determine a set of prescribed dispatch rules to dynamically reschedule existing patrol plans. The proposed solution not only reduces computational time required for training, but also improves the solution quality in comparison to an existing RL-based approach that relies on heuristic solver.

9199 Analysis of Heuristics for Vector Scheduling and Vector Bin Packing

Lars Nagel, Nikolay Popov, Tim Suess and Ze Wang

Fundamental problems in operational research are vector scheduling and vector bin packing where a set of vectors or items must be packed into a fixed set of bins or a minimum number of bins such that, in each bin, the sum of the vectors does not exceed the bin's vector capacity. They have many applications such as scheduling virtual machines in compute clouds where the virtual and physical machines can be regarded as items and bins, respectively. As vector scheduling and vector bin packing are NP-hard, no efficient exact algorithms are known. In this paper we introduce new heuristics and provide the first extensive evaluation of heuristics and algorithms for vector scheduling and bin packing including several heuristics from the literature. The new heuristics are a local search algorithm, a game-theoretic approach and a best-fit heuristic. Our experiments show a general trade-off between running time and packing quality. The new local search algorithm outperforms almost all other heuristics while maintaining a reasonable running time.

9237 Metaheuristics for optimal localisation of weather radars network over an airspace

Abdessamed Mogtit, Redouane Boudjemaa and Mohand Lagha

The objective of this study was to introduce a new approach, called "Co-Active Neural Network and Genetic Algorithms (CM-NNGA), which is designed for enhancing the coverage of a mixed weather radar network (WRN) in a predetermined spatial area. The model that we propose is intended for eliminating or minimizing the partial beam blockage (PBB) of each weather radar over a complex area, considering several constraints, namely the ground elevation, the radar beam elevation and the distance between two pairs of radars. The skipping technique has also been applied to decrease the computation time. The results reveal that the proposed model performs superior to well-known alternatives in quality with a somewhat higher computational time consumption, especially for large areas.

9680 Unleashing the potential of restart by detecting the search stagnation

Yoichiro Iida, Tomohiro Sonobe and Mary Inaba

SAT solvers are widely used to solve industrial problems owing to their exceptional performance. One critical aspect of SAT solvers is the implementation of restarts, which aims to enhance performance by diversifying the search. However, it is uncertain whether restarts effectively lead to search diversification. We propose to adapt search similarity index (SSI), a metric designed to quantify the similarity between search processes, to evaluate the impact of restarts. Our experimental findings, which employ SSI, reveal how the impact of restarts varies with respect to the number of restarts, instance categories, and employed restart strategies. In light of these observations, we present a new restart strategy called Break-out Stagnation Restart (BroSt Restart), inspired by a financial market trading technique. This approach identifies stagnant search processes and diversifies the search by shuffling the decision order to leave the stagnant search. The evaluation results demonstrate that BroSt Restart improves the performance of a sequential SAT solver, solving 19 more instances (+3%) than state-of-the-art solvers.

9999 Learning Heuristics for Combinatorial Optimization Problems with Deep Neural Networks

André Hottung

Solving real-world combinatorial optimization problems with traditional operations research methods can be a costly and time-consuming endeavor, often requiring the development of completely new methods or significant modification of existing techniques. In this talk, we explore the application of deep reinforcement learning to automate the development of problem-specific solution approaches. Rather than focusing on end-to-end solution generation, we investigate the use of machine learning to learn heuristic components for extensive search procedures. By automating the design of these components, the overall solution approach can be easily customized to the characteristics of specific problem instances, making optimization technologies more accessible to a wide range of organizations.