A multi-armed Bandit Approach for house ads recommendations

Marcel Goic; Nicolás Aramayo; Mario Schiappacasse

December 13, 2022 (Tuesday), 12:30 - Room 1209

Nowadays, websites use a variety of recommendation systems to decide the content to display to their visitors. In this work, we use a multi-armed bandit approach to dynamically select the combination of house ads to exhibit to a heterogeneous set of customers visiting the website of a large retailer. House ads correspond to promotional information displayed on the website to highlight some specific products and are an important marketing tool for online retailers. As the number of clicks they receive not only depends on their own attractiveness, but also on how attractive are other products displayed around them, we decide about complete collections of ads that capture those interactions. Moreover, as ads can wear out, in our recommendations we allow for non-stationary rewards. Furthermore, considering the sparsity of customer-level information, we embed a deep neural network to provide personalized recommendations within a bandit scheme. We tested our methods in controlled experiments where we compared them against decisions made by an experienced team of managers and the recommendations of a variety of other bandit policies. Our results show a more active exploration of the decision space and a significant increment in click-through and add-to-cart rates.

Keywords: Multi-Armed Bandits; House Ads; Personalization; Deep Learning

---

One-to-one and many-to-one Counterfactual Explanations for score-based models

Jasone Ramírez-Ayerbe; Emilio Carrizosa; Dolores Romero Morales

December 13, 2022 (Tuesday), 12:30 - Room 1209

Due to the increasing use of complex Machine Learning models in high-stakes decisions, it has become increasingly crucial to be able to understand how these models arrive at decisions. Assuming an already trained Supervised Classification model, an effective class of post-hoc explanations are counterfactual explanations, i.e., a set of changes that can be done to an instance such that the given Machine Learning model would have classified it in a different class. In this talk, for score-based classification models, we propose a novel Mathematical Optimization formulation for constructing different types of counterfactual explanations. We are able to generate the so-called collective counterfactual explanations, i.e., explanations for a group of instances in which we minimize the perturbation in the data (at the individual and group level) to have them labelled by the classifier in a given class. Moreover, we calculate representative counterfactuals for a batch of individuals, minimizing the distance while maximizing the probability of belonging to a specific class. Our methodology can generate sparse and plausible counterfactuals when dealing either with tabular data or with functional one. We illustrate our method with various real-world datasets.

Keywords: Counterfactual Explanations; Explainability; Mathematical Optimization; Machine Learning; Score-based Classification

---

MIP-Heuristics for the multi-level capacitated lot sizing problem

Fernanda Ueno; Marcos Furlan; Maristela Santos

December 13, 2022 (Tuesday), 12:30 - Room 1301

The multi-level lot-sizing problem determines a production plan to meet the demand for final items and/or their components without backlogging and considering the available production capacity. The objective of the problem is to minimize the setup and holding costs. To deal with this problem, we developed a fix-and-optimize heuristic. We investigate the impact on the quality of the solutions by considering classical partitions of the binary variables by period, items, and by resources. Also, we consider partitions of these variables constructed by unsupervised learning based on a work in the recent literature. This study uses the k-medoids to obtain non-balanced and balanced groups of the binary variables in our local search approach. The local search heuristic tends to local optima, so several strategies are used to build and use the obtained partitions. We use a relax-and-fix heuristic with decomposition by period and an LP-and-fix strategy to get feasible initial solutions. Solution

methods are implemented in Python integrated with the commercial LP/MIP solver Gurobi. Computational tests are performed using data from the literature. The solutions are compared with the BFO (bees-and-fix-and-optimize) heuristic, an approach of the literature developed to solve the problem that obtained the best results.

Keywords: Lot sizing problem; MIP-Based Heuristics; Fix-and-optimize heuristic; unsupervised learning

---

The economic lot sizing problem with remanufacturing and lost sales : complexity analysis and algorithms

Lucas Gana Reyes; Franco Quezada; Sebastián Dávila

December 13, 2022 (Tuesday), 12:30 - Room 1301

Motivated by global warming and the need of industrial companies to reduce both CO2 emissions and natural resource consumption, we study a new extension of the single-item lot-sizing problem first investigated by Wagner and Whitin (1958). In this extension, we consider the case of remanufacturing end-of-life products to obtain new products to meet demand over a finite discrete-time planning horizon. As in the classical lot-sizing problem, in case that in the period is produced a amount positive, then is incurred in a fixed cost -called setup cost-, besides of production cost per unit produced, and inventory holding cost per unit held in stock between two consecutive periods. In addition, due to a limited quantity of end-of-life products being returned by customers at each time period, the production system might not be able to satisfy the demand the customers demand on time. The corresponding demand is lost in this situation, incurring a high penalty cost to account for the loss of customer goodwill. Regarding that, the objective of this work is to build a production plan over a remanufacturing system so that the customers' demand is met in each time period and the total costs is minimized, i.e., the sum of setup, production, inventory holding costs, and lost sales, are minimized over the whole planning horizon. We thus first prove that the problem is NP-hard in its general case. Then, to solve the problem, structural and dominance properties of the optimal solution are investigated under several assumptions. These properties are then embedded within a branch-and-cut algorithm as valid inequalities to solve the problem. Computational experiments are carried out to assess its performance, and numerical results suggest a good performance of the proposed approach at solving the problem as compared to the one of a stand-alone mathematical programming solver.

Keywords: Lot sizing; Remanufacturing; Complexity; Valid inequalities; Branch and Cut algorithm

---

The economic lot-sizing problem with remanufacturing, backlogging and carbon emission constraint

Andrés Vallecilla; Franco Quezada; Sebastián Dávila

December 13, 2022 (Tuesday), 12:30 - Room 1301

The global warming makes us pay special attention to the greenhouse gas emissions produced by industrial companies, which significantly contribute to their release into the atmosphere that directly harms the care of the environment and the population's health. One way of mitigating this impact is to produce from end-of-life products rather than raw materials, reducing thus pollution and natural resources consumption. Governments have also committed to reducing their carbon emissions by promoting new regulations that restrict the production by industrial companies, which have responded by adapting the production processes to assuage these new limitations. We thus address a new extension of the single-item economic lot-sizing problem with carbon emission constraints. This extension aims to plan remanufacturing activities, which consist of deciding the timing and level of production as well as the resources to be used to meet the customers' demand in the most efficient and economical possible way. Moreover, in contrast to the classical manufacturing system, the remanufacturing production is limited by the number of end-of-life products returned by customers at each time period and, hence, the production system might not be able to deliver to all customers on time in a given period. We thus consider the case where it is still possible to satisfy the demand later than required (back-orders), incurring a penalty cost for delays proportional to the amount backlogged and the duration of the backlog. To address this problem, we study the structural