University of Antwerp

University of Antwerp
Operations Research Group
ANT/OR

ORBEL
THE BELGIAN OPERATIONAL RESEARCH SOCIETY

ORBEL 38

Antwerp, 8–9 February 2024

Booklet of Abstracts

Published by ORBEL, The Belgian Operations Research Society

https://orbel.be/

February 2024

Contents

Conference Information

1. Welcome to Antwerp!

ORBEL38 is the 38th conference of ORBEL, the Belgian Operational Research (OR) Society, Member of EURO, the association of European OR Societies, and Belgian representative of IFORS (International Federation of OR Societies).

The conference is intended as a meeting place for researchers, users and potential users of Operational Research, Statistics, Computer Science and related fields. It will provide managers, practitioners, and researchers with a unique opportunity to exchange information on quantitative techniques for decision making.

This year's meeting (ORBEL38) will take place in the city centre of Antwerp at the Faculty of Business and Economics of the University of Antwerp, and will be organized by ANT/OR - the Antwerp Operations Research Group.

2. ORBEL

ORBEL is the Belgian Society for the promotion of operational research and analytics. Its goal is to contribute to the science, the development, and the adoption of advanced tools, methods, and knowledge in these fields, by stimulating multidisciplinary research and scientific collaboration among its members, as well as fostering cooperation between academia, industry, and society in general, both on the national and the international level. Its main fields of interest include optimization, simulation, data science, and other quantitative methods for effective decision making in complex environments. ORBEL is a member of EURO, the Association of European Operational Research Societies, and IFORS, the International Federation of Operational Research Societies.

ORBEL / THE BELGIAN OPERATIONAL RESEARCH SOCIETY

3. Committees

3.1 Organizing Committee

- Christof Defryn
- Kenneth Sörensen
- Aline Janssens
- Pieter Deleye
- Gül Gündüz Mengübaş
- Ine Weyts

3.2 Scientific Committee

- Aghezzaf, El-Houssaine
- Arda, Yasemin
- Beliën, Jeroen
- Benoit, Dries
- Blondeel, Wouter
- Braekers, Kris
- Caris, An
- Chevalier, Philippe
- Coussement, Kristof
- Crama, Yves

- De Baets, Bernard
- De Causmaecker, Patrick
- Defryn, Christof
- De Smet, Yves
- Fortz, Bernard
- Goossens, Dries
- Janssens, Gerrit
- Kunsch, Pierre
- Labbé, Martine
- Leus, Roel
- Mélot, Hadrien
- Papadimitriou, Dimitri
- Paquay, Célia
- Pironet, Thierry
- Sartenaer, Annick
- Schauss, Pierre
- Schyns, Michaël
- Spieksma, Frits
- Sörensen, Kenneth
- Tancrez, Jean-Sébastien
- Van Utterbeeck, Filip
- Vandaele, Arnaud
- Vanden Berghe, Greet
- Vansteenwegen, Pieter
- Verbeke, Wouter
- Wauters, Tony
- Wittevrongel, Sabine

4. Practical Information

4.1 Presenting

Each presentation is scheduled for 20 minutes, including questions and discussion. Please adhere to this time frame.

Rooms are equipped with a computer attached to a beamer. Thanks in advance for using the provided computers as much as possible.

To ensure a timely progression of the session, we ask speakers to put their presentation files on the computers before the start of the session.

4.2 Chairing

Session chairs are listed in the detailed program. The session chair should briefly introduce the session topic as well as each of the speakers.

The chair should also remind the speakers to round off their talk when their time is almost up. If necessary, the chair may interrupt a speaker that is exceeding their alotted time slot.

We do not encourage switching rooms between talks in a session. In case of no shows, the chair may distribute the extra time between the speakers.

4.3 Venue

The conference is organized at the inner city campus of the University of Antwerp, building S.C. A map of the campus can be found below. You will enter the campus via Hof Van Liere, Prinsstraat 13, 2000 Antwerp.

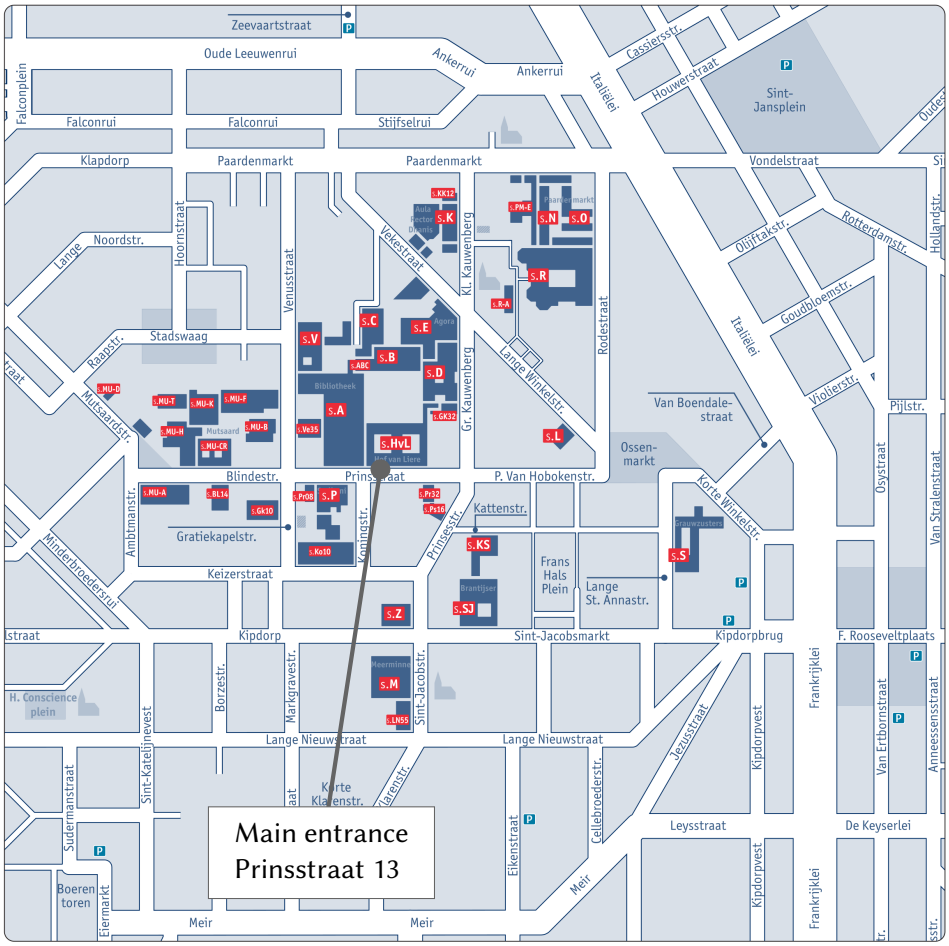

4.4 Conference dinner

The conference dinner will take place in the HORTA - Art Nouveau hall, just a short walk from the conference location. We are expected to take our seats at 19.30h.

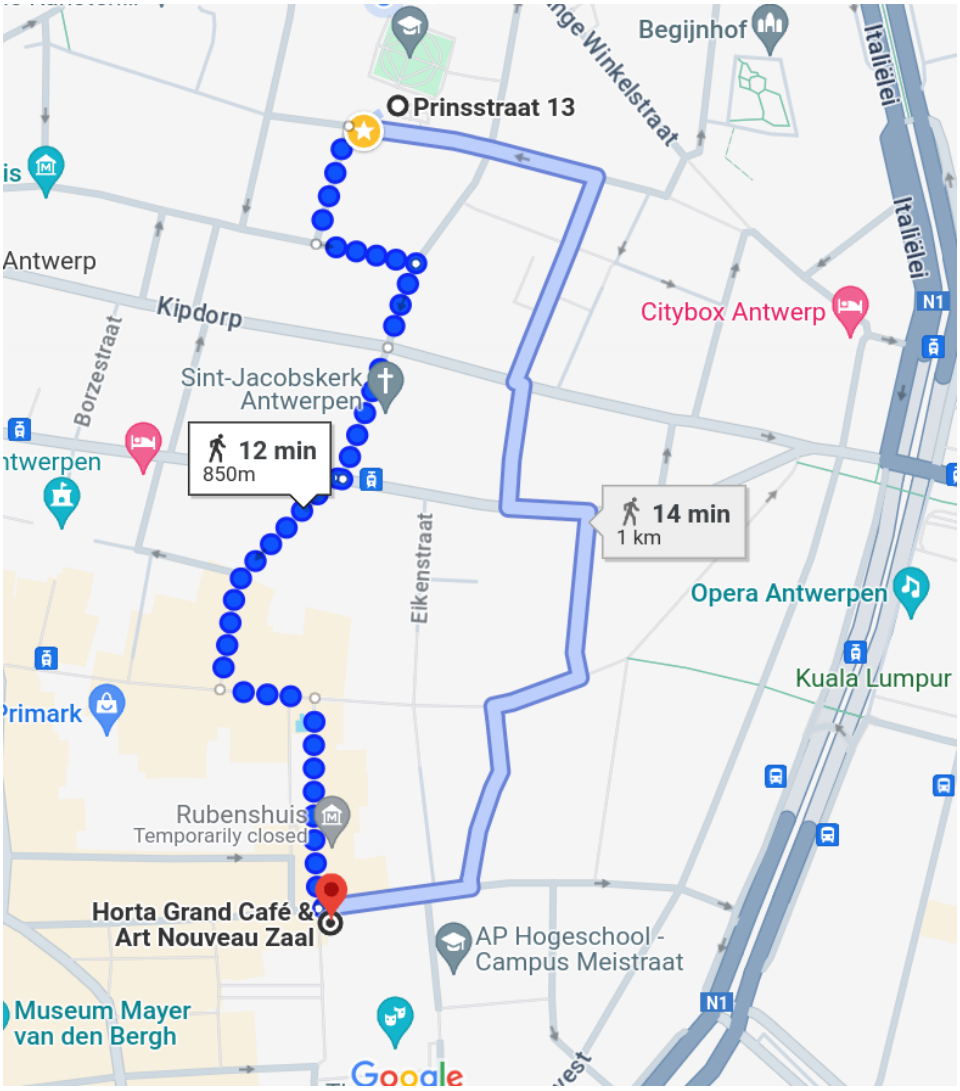

5. Sponsors

Universiteit Antwerpen

OMP.

ENGIE Impact

hexaly

6. Schedule at a glance

Thursday, 8 February 2024

08:45 - 09:30	Registration and welcome coffee	Corridors C
09:30 - 10:45	Opening session & Keynote Romero Morales	S.C.003
10:45 - 11:15	Coffee break	Corridors C
11:15 - 12:15	Parallel sessions TA	S.C.001, 002, 101, 102
12:15 - 13:30	Lunch	Corridors C
	ORBEL board meeting	Room “Lessius” (Building E)
13:30 - 14:50	Parallel sessions TB	S.C.001, 002, 101, 102
14:50 - 15:20	Coffee break	Corridors C
15:20 - 17:00	Parallel sessions TC	S.C.001, 002, 101, 102
17:00 - 18:00	ORBEL general assembly	S.C.002
19:30 - 22:30	Conference dinner	Horta

Friday, 9 February 2024

09:00 - 09:30	Welcome coffee	Corridors C
09:30 - 10:50	Parallel sessions FA	S.C.001, 002, 101, 102
11:00 - 12:00	Parallel sessions FB	S.C.001, 002, 101, 102
12:00 - 13:00	Lunch	Corridors C
13:00 - 14:20	Parallel sessions FC	S.C.001, 002, 101, 102
14:20 - 14:45	Coffee break	Corridors C
14:45 - 15:45	Keynote Gansterer	S.C.003
15:45 - 16:15	ORBEL award and closing session	S.C.003

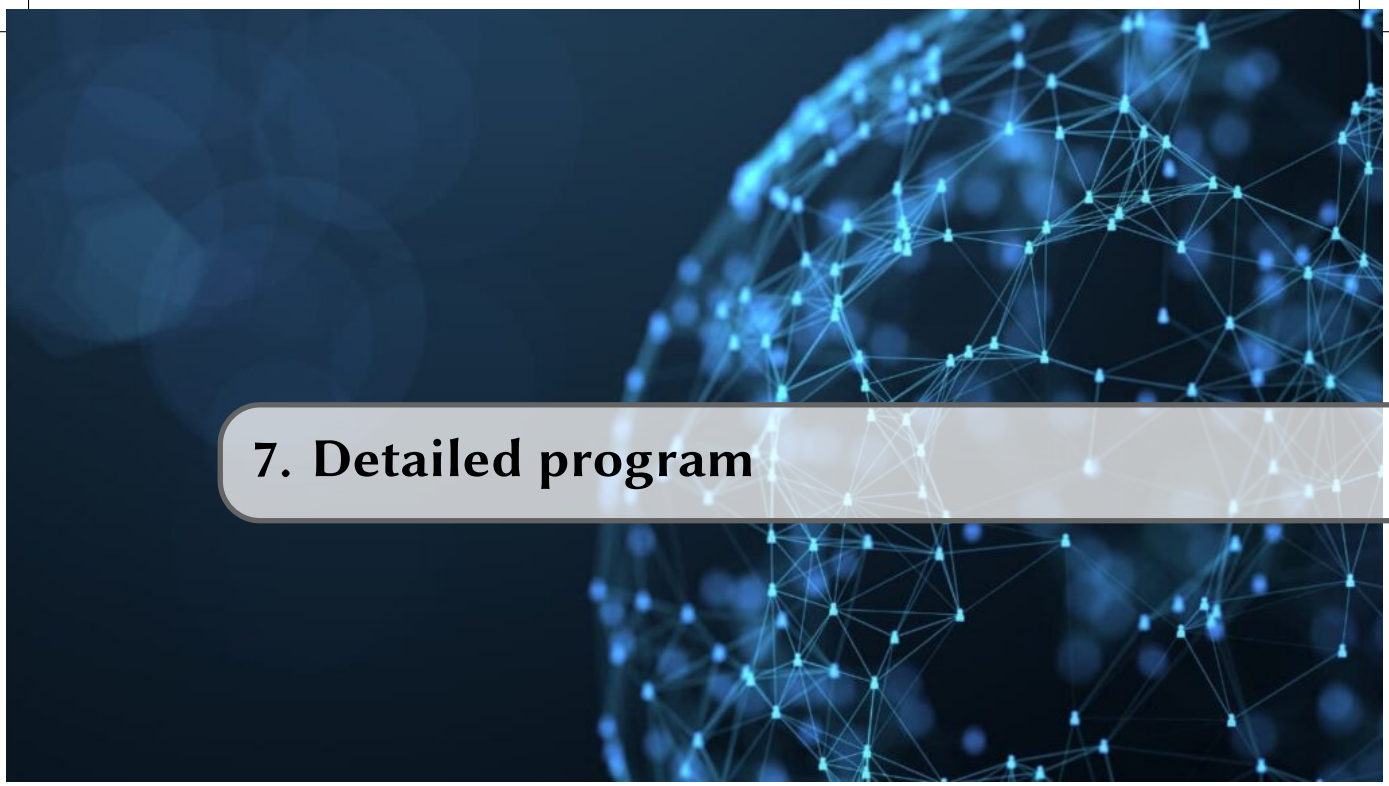

7. Detailed program

7.1 Thursday, 8 February 2024

08:45 - 09:30 **Registration and welcome coffee**
Corridors of building C

09:30 - 10:45 **Opening session**
Room: S.C.003

Welcome by Christof Defryn

Keynote Dolores Romero Morales
A tour through explainable and fair machine learning with an OR lens
Session chair: An Caris

10:45 - 11:15 **Coffee break**
Corridors of building C

11:15 - 12:15 **Parallel session TA1 - Forecasting**
Session chair: J. Van Belle
Room: S.C.001

11:15 - 12:15 **Parallel session TA2 - Packing**
Session chair: B. Zondervan
Room: S.C.002

11:15 - 12:15 **Parallel session TA3 - Vehicle routing**
Session chair: T. Iswari
Room: S.C.101

11:15 - 12:15 **Parallel session TA4 - Warehousing**
Session chair: A. Leroy
Room: S.C.102

12:15 - 13:30 **Lunch**
Corridors of building C
ORBEL board meeting
Room: Lessius

13:30 - 14:50 **Parallel session TB1 - Language Processing**
Session chair: M. Reusens
Room: S.C.001

13:30 - 14:50 **Parallel session TB2 - Timetabling**
Session chair: D. Goossens
Room: S.C.002

13:30 - 14:50 **Parallel session TB3 - Optimization**
Session chair: W. Vanroose
Room: S.C.101

13:30 - 14:50 **Parallel session TB4 - Assembly lines and manufacturing**
Session chair: H. Verplancke
Room: S.C.102

14:50 - 15:20 **Coffee break**
Corridors of building C

15:20 - 17:00 **Parallel session TC1 - Machine learning**
Session chair: A. Stevens
Room: S.C.001

15:20 - 17:00 **Parallel session TC2 - Matheuristics**
Session chair: F. Phillipson
Room: S.C.002

15:20 - 17:00 **Parallel session TC3 - Graph theory**
Session chair: T. Vancaeyzeele
Room: S.C.101

15:20 - 17:00 **Parallel session TC4 - Algorithmic design**
Session chair: K. Sörensen
Room: S.C.102

17:00 - 18:00 **ORBEL general assembly**
Room: S.C.002

19:30 **Conference dinner**
Horta - Art Nouveau Hall

7.2 Friday, 9 February 2024

09:00 - 09:30 **Welcome coffee**
Corridors of building C

09:30 - 10:50 **Parallel session FA1 - Machine learning**
Session chair: J. Sanchez Ramirez
Room: S.C.001

09:30 - 10:50 **Parallel session FA2 - Economics and game theory**
Session chair: B. Lin
Room: S.C.002

09:30 - 10:50 **Parallel session FA3 - Collaboration and sharing**
Session chair: M. Wens
Room: S.C.101

09:30 - 10:50 **Parallel session FA4 - Vehicle routing**
Session chair: A. Maharani
Room: S.C.102

11:00 - 12:00 **Parallel session FB1 - Machine learning**
Session chair: A. Thuy
Room: S.C.001

11:00 - 12:00 **Parallel session FB2 - Energy**
Session chair: J. Springael
Room: S.C.002

11:00 - 12:00 **Parallel session FB3 - Graph theory**
Session chair: B. De Clerck
Room: S.C.101

11:00 - 12:00 **Parallel session FB4 - Warehousing**
Session chair: B. Raa
Room: S.C.102

12:00 - 13:00 **Lunch**
Corridors of building C

13:00 - 14:20 **Parallel session FC1 - Feature selection and parameter tuning**
Session chair: B. Coquelet
Room: S.C.001

13:00 - 14:20 **Parallel session FC2 - Scheduling**
Session chair: Leonie Gallois
Room: S.C.002

13:00 - 14:20 **Parallel session FC3 - Orbel Award Nominees**
Session chair: R. Leus
Room: S.C.101

Development of a Variable Neighborhood Search algorithm for scheduling and routing fully flexible 'on-demand-only' feeder bus systems
A. Hens

Optimization method to integrate driver consistency and route balancing in a heterogeneous multi-period dial-a-ride problem
M. Hubert

Inventory Management in a Synchromodal Transport Network - An Inventory Replenishment Heuristic with Order Expediting
F. Sneyers

13:00 - 14:20 **Parallel session FC4 - Vehicle routing**
Session chair: L. Melis
Room: S.C.102

Systematic Derivation of a Global Constraint for Routing Optimization
R. De Landtsheer 198

The Dial-a-Ride Problem with Limited Pickups per Trip
B. Zhao, K. Wang, W. Wei, and R. Leus 201

Delivering efficiency: The parcel locker puzzle
L. Melis, A. Ranjbari, and V. V. Gayah 203

14:20 - 14:45 **Coffee break**
Corridors of building C

14:45 - 15:45 **Plenary session - Keynote Margaretha Gansterer**
Last-mile delivery with shared resources: modeling and recent developments
Session chair: Kenneth Sörensen
Room: S.C.003

15:45 - 16:15 **ORBEL award and closing session**
Room: S.C.003

8. Plenary talks

8.1 A tour through Explainable and Fair Machine Learning with an OR lens

Dolores Romero Morales, Copenhagen Business School, Denmark

The use of Artificial Intelligence and Machine Learning algorithms is ubiquitous in Data Driven Decision Making. Despite their excellent accuracy, these algorithms are often criticised for their lack of transparency. Algorithms such as Random Forest, XGBoost and Deep Learning are often seen as black boxes, for which it is difficult to explain their predictions. In addition, when applied to sensitive situations with consequential impact on citizens' lives, including access to social services, lending decisions or parole applications, this opaqueness may hide unfair outcomes for risk groups. Therefore, there is an urgent need to strike a balance between three goals, namely, accuracy, explainability and fairness. In this presentation, and with an Operations Research lens, we will navigate through some novel Machine Learning models that embed explainability and fairness in their training.

Dolores Romero Morales is a Professor in Operations Research at Copenhagen Business School. Her areas of expertise include Data Science, Supply Chain Optimization and Revenue Management. In Data Science she investigates explainability/interpretability, fairness and visualization matters. In Supply Chain Optimization she works on environmental issues and robustness. In Revenue Management she works on large-scale network models. Her work has appeared in a variety of leading scholarly journals, including European Journal of Operational Research, Management Science, Mathematical Programming and Operations Research, and has received various distinctions. Currently, she is Editor-in-Chief to TOP, the Operations Research journal of the Spanish Society of Statistics and Operations Research, and an Associate Editor of Journal of the Operational Research Society, Omega and the INFORMS Journal on Data Science. She has worked with and advised various companies on these topics, including IBM, SAS, KLM and Radisson Edwardian Hotels, as a result of which these companies managed to improve some of their practices. SAS named her an Honorary SAS Fellow and member of the SAS Academic Advisory Board. She currently leads the EU H2020-MSCA-RISE NeEDS project, which has a total of 15 participants and a budget of more than €1.000.000 for intersectoral and international mobility, with the aim to improve the state of the art in Data Driven Decision Making. Dolores joined Copenhagen Business School in 2014. Prior to coming to Copenhagen Business School, she was a Full Professor at University of Oxford (2003-2014) and an Assistant Professor at Maastricht University (2000-2003). She has a BSc and an MSc in Mathematics from Universidad de Sevilla and a PhD in Operations Research from Erasmus University Rotterdam.

8.2 Last-mile delivery with shared resources: modeling and recent developments

Margaretha Gansterer, University of Klagenfurt, Austria

The Sharing Economy is on the rise. Traditional business models have to be adapted and players have to succeed in a world of shared idle capacities on digital platforms. Innovative concepts related to shared transportation resources include collaborative vehicle routing or crowd delivery systems. A collaboration can be described as a partnership between two or more companies to optimize operations by making joint decisions and sharing information, resources, or profits. Crowdsourced deliveries are either conducted by freelance drivers or by occasional drivers, where the former conduct several tasks over a longer time horizon typically being active on different platforms, while the latter offer their service on a more irregular basis. Both of these concepts have their advantages but also specific challenges. In case of horizontal transport collaborations we distinguish between centralized and decentralized settings. In the former, it is assumed that one fully informed decision maker exists, while in the latter, decisions under incomplete information have to be orchestrated. Crowd delivery concepts, however, highly rely on the willingness of external drivers to offer their service. Hence, attractive offers have to made. We focus on the generation of bundles of requests such that the distribution of orders assigned to external drivers and to a company-owned fleet is optimized. The talk will cover modeling approaches, recent developments, and promising future directions.

Margaretha Gansterer currently holds a full professor position at the University of Klagenfurt. Since October 2019 she is serving as Head of the Department of Operations, Energy, and Environmental Management and since January 2022 as Dean of the Faculty of Management, Economics & Law. Margaretha obtained her PhD and her Habilitation at the University of Vienna. She also held research positions at Johannes Kepler University Linz and the Austrian Research Centers, and a professor position at Otto-von- Guericke University Magdeburg. From 2020 to 2021 she served as International Liaison (Europe, Middle East, Africa) of the INFORMS Transportation Science and Logistics (TSL) society. She is currently serving as president-elect and will take over the TSL presidency in 2024. Her research interests lie in the application of quantitative methods to planning problems arising in operations management and logistics with a special focus on collaborative operational planning. She has published more than 50 refereed papers in international journals, books, and conference proceedings. Margaretha has served as a reviewer for numerous international journals such as Transportation Science, Transportation Research Part B, and the European Journal of Operational Research. Since 2015 she is leading research projects financed by the Austrian Science Fund (FWF).

Abstracts

Towards sustainable power systems: improving battery degradation forecasts

E. Grégoire
Vrije Universiteit Brussel, Data Analytics Laboratory
emilie.gregoire@vub.be

S. Verboven
Vrije Universiteit Brussel, Data Analytics Laboratory
sam.verboven@vub.be

Keywords: Sustainability Time Series Forecasting Multi-Task Learning.

Abstract

The task of developing a reliable and long-lasting battery system remains a major obstacle to the expansion of the electric vehicle fleet and the improvement of storage systems in renewable energy plants. To overcome this challenge, improving battery ageing models to accurately forecast the battery degradation trajectory is a complex but critical step. In this work, we introduce a novel sequence-to-sequence multi-task learning (MTL) method. Empirical results on battery ageing datasets show that our model enhances the data efficiency and precision of capacity and power degradation forecasts for lithium-ion battery cells. Furthermore, our model generates new insights into the trade-offs and key decisions for incorporating battery ageing models into the MTL paradigm.

1 Problem setting

In order to enhance the positive impact of lithium-ion batteries on a carbon-neutral future, significant improvements in battery performance and reliability are needed [1]. To address these requirements, a crucial step is to enable accurate forecasting of lithium-ion battery degradation, which is reflected by capacity loss and an increase in the internal resistance. Recently, multi-task learning (MTL) has emerged as a promising approach to improve the robustness and generalization performance of data-driven forecasting methods [2].

In MTL, the model parameters are partially shared among the different prediction tasks. These parameters are updated through gradient descent by using the signal of the composite loss $\mathcal{L}_{Tot}$, which corresponds to the weighted sum of the task-specific capacity

and resistance losses, denoted by $\mathcal{L}_C$ and $\mathcal{L}_R$ respectively,

$$\mathcal{L}_{Tot} = w_C \mathcal{L}_C + w_R \mathcal{L}_R. \tag{1}$$

The coefficients w_C and w_R are the corresponding task-specific weights. In [2], the benefits of a first instantiation of an MTL network were demonstrated. In this work we explore the design space of MTL networks for battery degradation forecasting, and quantify the impact of the corresponding choices.

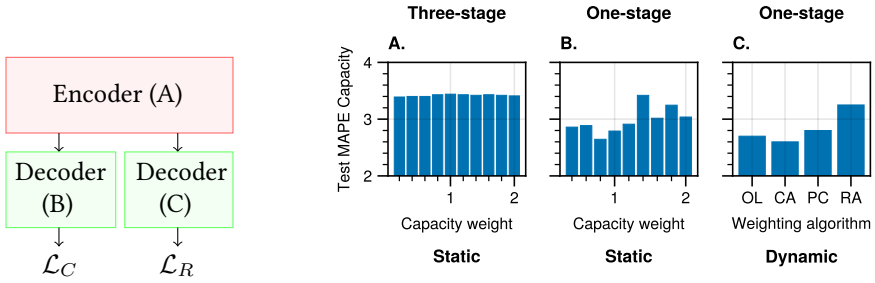

(a) Encoder-decoders

(b) Test MAPE for different configurations

Figure 1: Left Figure (a): the encoder (A) and decoder (B and C) modules. Right Figure (b): bar plots representing the mean absolute percentage errors (MAPE) for the capacity predictions. The comparison between panels **A** and **B** illustrates the difference between the three-stage and one-stage setups (i) and the corresponding impact of the capacity task weight. Next, panels **B** and **C** demonstrate the difference between static and dynamic task weighting (ii) for several weighting algorithms.

2 Methodology and Results

We explore the space of choices concerning (i) how to combine several (task-specific) modules (Figure 1a) and (ii) how to assign weights w_i to the task losses $\mathcal{L}_i$ (equation 1). To address (i) we compare the performance of two setups: the one-stage setup where all modules are trained simultaneously (ABC) and the three-stage setup where different subsets of the modules (AB-AC-ABC) are trained during each stage. Regarding (ii), we analyze the impact of various static weights and compare the performance of state-of-the-art dynamic weighting algorithms. We show that distinct choices for (i) and (ii) yield significant differences (Figure 1b) in generalization performance and should be optimized to obtain better results. The experimental evaluation is conducted on two open-source datasets. The encoder and decoder modules consist of Bidirectional Long Short-Term Memory (LSTM) layers.

3 References

[1] Jones, Penelope K. and Stimming, Ulrich and Lee, Alpha A. (2022). Impedance-based forecasting of lithium-ion battery performance amid uneven usage. Nature Energy.

[2] Li, Weihan and Zhang, Haotian and van Vlijmen, Bruis and Dechent, Philipp and Sauer, Dirk Uwe. (2022). Forecasting battery capacity and power degradation with multi-task learning. Energy Storage Materials. (vol 53, 453-466)

An Exploration of Exact Methods for Effective Network Failure Detection and Diagnosis

A. Burlats C. Pelsser P. Schaus
UCLouvain, ICTEAM
firstname.surname@uclouvain.be

Keywords: Boolean tomography, Network supervision.

Computer networks form the backbone of modern digital communication, and their reliability is crucial for maintaining seamless connectivity across various sectors. Failures within these networks can have significant consequences, leading to service disruption and potential financial loss. As such, it is essential to develop efficient and accurate methods for detecting and diagnosing network failures, enabling swift recovery and minimizing the impact on end-users.

In this study, we focus on Boolean network tomography, a research field that holds great promise for enhancing the resilience of networks. Boolean network tomography combines end-to-end measures with inference algorithms to estimate the state of different elements in the network. Its advantage is that it only requires a subset of nodes to be monitors and supervise an entire network.

With this approach, monitors send messages to each other through *measurement paths*. When a failure occurs on a node, all paths that cross it fail. Thus, the failure can be detected by observing if some measurement paths are not working. If the set of failed measurement paths forms a unique signature, then it is even possible to identify the failed node.

Our investigation focuses on minimizing the number of designated monitor nodes while ensuring some level of quality of network monitoring. This is crucial for minimizing monitoring costs without compromising the network's overall health and performance. We explore two critical monitoring challenges: the *cover problem*, which seeks to detect failures, and the *1-identifiability problem*, which requires pinpointing the exact failing node.

Conceptually, a node failure in a network results in the disruption of all paths traversing it. These affected paths collectively constitute the *symptom* associated with the failing node. A network is covered if there is a non-empty symptom for each node. Additionally, a network is considered 1-identifiable if every node possesses a unique, non-empty symptom, thereby serving as an identifier for the node in the event of a failure. By compiling a comprehensive list of these identifiers, one can efficiently diagnose a failure by simply observing the disrupted paths and cross-referencing a precomputed table that maps the

failed paths (symptoms) to the corresponding node.

An important assumption of the considered networks in this study is that the routes between any pairs of nodes are imposed by the routing protocol and known by the planning tool that will select the monitors. A pair of monitors is only able to verify the status of those routes. In practice network operators usually configure link (IGP) weights to influence where the traffic flows in the network assuming they follow shortest paths (see for instance [1] for optimizing IGP weights). Alternatively, other protocols such as segment routing or MPLS [2, 3, 4] make it possible to introduce deviations or explicit route setups between pairs of nodes, deviating from shortest paths. For all these protocols, the monitors are able to determine which data paths between them are affected by a failure.

We introduce and compare an integer linear programming, a constraint programming and a MaxSAT model to solve both the node cover and 1-identifiability NP-hard problems. To reduce the search space, we propose redundant constraints and problem reductions exploiting the network topologies. Using 625 real network topologies, we compare the solutions returned by each model to a specialized version of a greedy algorithm called MNMP [5], tailored for 1-identifiability. Our findings reveal that the introduced models reduce the amount of monitor nodes compared to the greedy approach for 35 instances of the cover problem and 157 instances of the 1-identifiability problem, paving the way for more robust and reliable telecommunication networks.

References

[1] Bernard Fortz and Mikkel Thorup. Internet traffic engineering by optimizing ospf weights. In *Proceedings IEEE INFOCOM 2000. conference on computer communications. Nineteenth annual joint conference of the IEEE computer and communications societies (Cat. No. 00CH37064)*, volume 2, pages 519–528. IEEE, 2000.

[2] Renaud Hartert, Pierre Schaus, Stefano Vissicchio, and Olivier Bonaventure. Solving segment routing problems with hybrid constraint programming techniques. In *Principles and Practice of Constraint Programming: 21st International Conference, CP 2015, Cork, Ireland, August 31–September 4, 2015, Proceedings 21*, pages 592–608. Springer, 2015.

[3] Renaud Hartert, Stefano Vissicchio, Pierre Schaus, Olivier Bonaventure, Clarence Filsfils, Thomas Telkamp, and Pierre Francois. A declarative and expressive approach to control forwarding paths in carrier-grade networks. *ACM SIGCOMM computer communication review*, 45(4):15–28, 2015.

[4] Youngseok Lee, Yongho Seok, Yanghee Choi, and Changhoon Kim. A constrained multipath traffic engineering scheme for mpls networks. In *2002 IEEE International Conference on Communications. Conference Proceedings. ICC 2002 (Cat. No. 02CH37333)*, volume 4, pages 2431–2436. IEEE, 2002.

[5] Liang Ma, Ting He, Ananthram Swami, Don Towsley, and Kin K. Leung. On optimal monitor placement for localizing node failures via network tomography. *Performance Evaluation*, 91:16–37, September 2015.

Forecasting Business Process Models

J. Van Belle
KU Leuven, LIRIS
jente.vanbelle@kuleuven.be

Jo. De Smedt
KU Leuven, LIRIS
johannes.desmedt@kuleuven.be

A. Polyvyanyy
The University of Melbourne
artem.polyvyanyy@unimelb.edu.au

J. De Weerdt
KU Leuven, LIRIS
jochen.deweerdt@kuleuven.be

Keywords: Predictive process monitoring, Hierarchical time series forecasting, Forecast combination

Business process management is a domain concerned with the identification, design, implementation, execution, and improvement of business processes that underpin the operations of an organization. Over time, there has been a shift away from classical qualitative approaches towards data-centric techniques, which led to the emergence of the process mining field.

The key element in process mining is the use of event log data produced by process-aware information systems. These event logs contain sequences of activities tied to a particular case or process instance, such as a loan application, where a client fills in an application form, goes through the required checks, and finally receives a formal response. Traditionally, process mining relied on backward-looking and reflective analysis. More recently, however, forward-looking or predictive approaches were developed in which event logs are used to predict the future development of cases, e.g., to predict whether a loan application will be accepted or the remaining time until a decision given the candidates current progress through their application and their data. These predictive process monitoring solutions typically focus on predictions for individual process instances or cases. To address this limitation, De Smedt et al. (2021) propose a paradigm shift towards predictive modeling also at the process model level to predict how the system as a whole will evolve over time.

In this paper, we propose a generic three-step methodology to robustly forecast entire process models from an event log of the corresponding historical process. The first step involves transforming the event log into a collection of directly-follows (DF) time series, with each DF time series capturing the evolution in the number of transitions between a specific pair of directly subsequent activities during a specified unit of time. The DF time series thus capture the control flow aspect of the historical process. Each DF time series can then be forecasted independently or jointly by relying on univariate or multivariate time series forecasting techniques, respectively. In this work, we focus on univariate

time series forecasting methods because they can generally be applied, irrespective of the characteristics of the event log. Since event log data is typically collected on a nearly continuous basis nowadays, the unit of time used to construct the DF time series can be specified by the user by simply using a different aggregation filter. Choosing a single data frequency while taking into account the desired forecast period length is, however, a non-trivial problem. Therefore, in the second step, we generate different views of each DF time series through non-overlapping temporal aggregation to construct temporal hierarchies. Non-overlapping temporal aggregation is known to strengthen or attenuate different elements of a time series. The different views of the data therefore allow us to derive different insights. Various forecasting methods are applied to produce several base forecasts for each view. For each forecasting method, we also create a reconciled version of the base forecasts, ensuring consistency across time, i.e., imposing that lower-level forecasts properly add up to higher-level forecasts (Athanasopoulos et al., 2017). The base forecasts and reconciled forecasts are then used in forecast pooling to generate a single consolidated forecast for each DF (Kourentzes et al., 2019). Finally, in the third step, we obtain a robust process model forecast by combining a filtered set of consolidated DF forecasts.

We apply this methodology to real-life event logs to illustrate how the forecasted process models provide a substantial benefit over process models obtained through a classical backward-looking process discovery approach. The generated process model forecasts are evaluated from both a time series forecasting and a process mining perspective.

References

Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., and Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, 262(1):60–74.

De Smedt, J., Yeshchenko, A., Polyvyanyy, A., De Weerdt, J., and Mendling, J. (2021). Process model forecasting using time series analysis of event sequence data. In *International Conference on Conceptual Modeling*, pages 47–61. Springer.

Kourentzes, N., Barrow, D., and Petropoulos, F. (2019). Another look at forecast selection and combination: Evidence from forecast pooling. *International Journal of Production Economics*, 209:226–235.

Decoupling Geometry from Optimization, an Open-Source Collision Detection Engine for 2D irregular Cutting and Packing problems

J. Gardeyn
KU Leuven, NUMA Belgium
jeroen.gardeyn@kuleuven.be

T. Wauters
KU Leuven, NUMA Belgium
tony.wauters@kuleuven.be

Keywords: cutting and packing, irregular, collision detection

2D cutting and packing (C&P) problems involve placing a set of smaller items within the bounds of a larger container. This broad class of problems can be very challenging to solve, especially when the items and containers are *irregular* (non-rectangular) in shape. This difficulty has led to a highly active research domain featuring many approaches to many variants of the problem.

While the objectives and particulars of these irregular C&P problems can differ, they all share the need for a specific feasibility check: determining whether or not an item can be placed at a certain position. This check is particularly complex due to the geometric irregularity of the items and containers. The general problem of *collision detection* also prominently arises in other fields, such as computer graphics and simulations. This seemingly simple problem is so challenging that entire books have been dedicated to this subject.

For C&P problems, there is currently no sufficiently general approach to tackle this task in an easy and efficient manner. This not only heightens the barrier to entry for those who must address such problems, but also results in different researchers having to continuously reinvent the wheel whenever they are addressing new problems or developing new approaches for existing problems.

This work aims to decouple geometry from optimization and develop a high-performing adaptable *engine*, capable of efficiently handing the geometric component of irregular C&P problems. We envisage two target audiences. First, there are those who simply want to focus on their optimization problem at hand and who would therefore greatly benefit from having an engine they can incorporate into their own methodology. Such an engine would essentially outsource the geometric challenge and enable them to focus their efforts on developing smart solution methods. The second target audience are those who, rather than solving problems themselves, might have good ideas concerning how to further improve and refine this open-source engine.

Solving a tri-objective menu planning problem within the SHARP framework.

F. Martos-Barrachina
Universidad de Malaga, Department of Statistics and Operational Research
fmeco@uma.es

L. Delgado-Antequera
Universidad de Malaga, Department of Applied Economics (Mathematics)

M. Hernandez Huelin
Universidad de Malaga, Department of Applied Economics (Mathematics)

Keywords: Combinatorial Optimization, Multiobjective Optimization, Menu Planning Problem

1 The MPP Problem.

The Menu Planning Problem (MPP) can be conceptualized as a basic combinatorial optimization problem in the field of Operations Research (OR). In its essence, the MPP involves determining an optimal combination of items (recipes) to include in a menu to meet certain criteria while satisfying constraints.

In the context of combinatorial optimization, the goal is to find the best possible arrangement or combination from a finite set of elements. In the MPP, these elements are the various items that could potentially constitute a menu. The optimization objective typically involves maximizing acceptability or minimizing cost.

Mathematically, the MPP is often formulated as an binary variable problem. The decision variables represent whether each menu item is included in a particular slot of the menu or not. The constraints capture the limitations and requirements imposed on the menu, particularly the nutritional and cultural ones.

2 The SHARP approach.

The SHARP approach in the Menu Planning Problem (MPP) is a comprehensive framework that integrates five critical criteria into the decision-making process. The acronym

SHARP stands for Sustainability, Health, Affordability, Reliability, and Palatability, encapsulating the key dimensions that are essential in crafting well-rounded and effective menus.

In the most basic approach, cost is THE objective, and usually health (nutritiousness) is encapsulated in the constraints. In these holistic approach, all five dimensions appear. We have incorporated all of the criteria into the problem, ending up with a tri-objective problem in which cost, palatability and sustainability are our goals.

3 Solving the multicriteria decision problem using an E-WAF.

we have employed the Extended Wierzbicki Achievement Function to facilitate decision-making by simultaneously considering three key objectives: Cost, Palatability, and Sustainability. The Extended Wierzbicki Achievement Function is an extension of the classical achievement scalarizing function used in multi-objective optimization.

The Extended Wierzbicki Achievement Function enables us to navigate the trade-offs between these three objectives, transforming them into a single aggregated measure that reflects the overall performance of a given menu. By doing so, this approach provides a practical and comprehensive method for decision-makers to evaluate and choose menus that strike an optimal balance between cost-effectiveness, palatability, and sustainability.

This extended function is particularly useful in the Menu Planning Problem because it acknowledges the inherently multi-faceted nature of the decision space. It allows us to address the complexity of optimizing menus, where various, sometimes conflicting, objectives must be considered simultaneously. The Extended Wierzbicki Achievement Function contributes to a more nuanced and balanced decision-making process, aligning with the holistic approach inherent in the SHARP framework.

4 References

[1] Martos-Barrachina, F., Delgado-Antequera, L., Hernandez, M., and Caballero, R. (2022) An extensive search algorithm to find feasible healthy menus for humans. Operational Research, 22(5):36, 2022.

[2] van Dooren, C., and Aiking, H. (2015). Defining a nutritionally healthy, environmentally friendly, and culturally acceptable low lands diet. The International Journal of Life Cycle Assessment, 21(5):688–700, nov 2015. doi: 10.1007/s11367-015-1007-3.

[3] Wierzbicki, A. (1982). A mathematical basis for satisficing decision making. Mathematical modelling, 3(5): 391–405.

Comparing mesh and voxel-based approaches for 3D irregular Cutting and Packing Problems

B. Zondervan, T. Wauters
KU Leuven, Department of Computer Science, CODeS research group
bart.zondervan@kuleuven.be, tony.wauters@kuleuven.be

Keywords: Voxelization, Nofit voxel

In logistics and manufacturing, a common task is to arrange objects within a three dimensional volume. This task presents a considerable challenge, particularly when dealing with irregularly shaped objects that can be rotated in multiple ways. These optimization problems are known as cutting and packing problems and are unique in the fact that they combine geometric computations with combinatorial optimization.

At the heart of cutting and packing problems lies the question of how to represent objects. Two popular techniques for representing objects are triangular meshes and voxelization. Triangular meshes represent the objects by a triangulated surface. Voxelization approximates the objects by 3-dimensional pixels called voxels.

Efficiently checking for overlaps between objects is crucial in cutting and packing problems. One method for this involves so-called nofit voxels. The nofit voxel for a pair of objects is the set of all voxels that one object can be placed such that it intersects the other object. By computing the nofit voxel in a preprocessing phase, it becomes possible to quickly decide if two object overlap, which is a question that occurs very often in cutting and packing problems.

Moreover, voxelization and nofit voxels offer the advantage of parallelization on GPUs, a crucial feature for speeding up computations. The main goal of this talk is to explain the pros and cons of both representation techniques and to show results regarding the fast parallel implementation of voxelization.

Leveraging Contextual Information for Robustness in Vehicle Routing Problems

A. İ. Mahmutoğulları
Department of Computer Science, KU Leuven
irfan.mahmutogullari@kuleuven.be

T. Guns
Department of Computer Science, KU Leuven
tias.guns@kuleuven.be

Keywords: Routing, Robust optimization, Quantile predictions, Contextual information, Time windows.

We investigate the benefit of using contextual information in data-driven demand predictions to solve the robust capacitated vehicle routing problem with time windows. Instead of estimating the demand distribution or its mean, we introduce contextual machine learning models that predict demand quantiles even when the number of historical observations for some or all customers is limited. We investigate the use of such predicted quantiles to make routing decisions, comparing deterministic with robust optimization models. Furthermore, we evaluate the efficiency and robustness of the decisions obtained, both using exact or heuristic methods to solve the optimization models. Our extensive computational experiments show that using a robust optimization model and predicting multiple quantiles is promising when substantial historical data is available. In scenarios with a limited demand history, using a deterministic model with just a single quantile exhibits greater potential. Interestingly, our results also indicate that the use of appropriate quantile demand values within a deterministic model results in solutions with robustness levels comparable to those of robust models. This is important because, in most applications, practitioners use deterministic models as the industry standard, even in an uncertain environment. Furthermore, as they present fewer computational challenges and require only a single demand value prediction, deterministic models paired with an appropriate machine learning model hold the potential for robust decision-making.

Acknowledgements This research received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant No. 101002802, CHAT-Opt).

Advancements in Space Logistics: The Moving Target TSP Paradigm

P. Deleye
University of Antwerp, Department of Engineering management
pieter.deleye@uantwerpen.be

C. Defryn
University of Antwerp, Department of Engineering management
christof.defryn@uantwerpen.be

Keywords: Travelling Salesman Problem; moving target routing; OR in space

The Traveling Salesman Problem (TSP) is a combinatorial optimization problem which aims to find a tour visiting a set of vertices, such that each vertex is visited exactly once (Hamiltonian cycle). A large number of TSP variations have been studied extensively within the operations research field, as these allow the mapping of several real-world problems. A special case of TSP is the one in which the vertices are not static, but their positions change over time. This variation is known as the Moving Target TSP.

The problem was first introduced by Helvig, Robins, and Zelikovsky (1998) within a single dimension, i.e, all points lie on a single line crossing the origin. They describe more generally the moving target TSP as a time-dependent generalisation of the TSP where a pursuer must intercept a set of targets which move at a constant velocity. Although a constant velocity was assumed for their targets, these could either move towards or away from the origin. The authors propose exact and approximate algorithms for their moving target TSP considering a small number of moving points (all on one line, with or without resupply, i.e., returning to the depot before continuing to the next target, and the case where the points never reach the origin.)

Over the last few decades, the problem has been enriched by several authors, each adding one (or some) of the following elements: increased number of moving targets (from 5 up to 100), different velocities of the moving targets (velocities are typically between 0 and 2 m/s; either constant for all targets or stochastic per target), the possibility to resupply during the trip, release times and deadlines for individual targets, and various trajectories (linear or parabolic).

Applications of the moving target TSP are often found in a military / defense context. Examples include the interception of multiple missiles to protect a high-value target, an aircraft hitting moving ground targets, the use of unmanned aerial vehicles to intercept

multiple target persons in a dispersing crowd, or resupplying a fleet of patrol boats. For more information we refer to the work of Uçar and Işleyen (2020). The moving target TSP also features specific characteristics that might be highly valuable for applications in space, such as the problem of servicing a constellation of satellites in orbit. To the best of our knowledge, such applications are often overlooked but could be highly relevant in today's rapidly changing space market.

Our presentation will feature our first insights in the academic literature on the moving target TSP with a focus on applications in space. As such, we aim to motivate the need for models and algorithms that can deal with cyclical and elliptical target itineraries as well as orbital dynamics.

References

[1] Helvig, C. S., Robins, G., & Zelikovsky, A. (1998). Moving-target TSP and related problems. In *European Symposium on Algorithms* (pp. 453-464). Berlin, Heidelberg: Springer Berlin Heidelberg.

[2] Uçar, U. & Işleyen, S.K. (2020). A Survey of Moving Target Traveling Salesman Problem. In book: *Human-Computer Interaction*, edited by Turgut Ozseven. Nova Science Publishers, Inc.

A study on the benefits of integrating inventory and routing decisions in a city logistics context

T. Iswari
Hasselt University, Research Group Logistics
titi.iswari@uhasselt.be

A. Caris
Hasselt University, Research Group Logistics
an.caris@uhasselt.be

K. Braekers
Hasselt University, Research Group Logistics
kris.braekers@uhasselt.be

Keywords: Inventory, Routing, City-logistics

This study investigates the potential of utilizing an Urban Consolidation Center (UCC) or a city hub as an intermediate storage facility within a B2B two-echelon urban distribution system. The main objective is to assess the advantages of integrating inventory and routing decisions in contrast to a traditional approach where the city hub functions only as a consolidation point. Two scenarios are explored in this study: a traditional one involving sequential inventory and routing decisions, and an integrated one where these decisions are made simultaneously.

A traditional problem in which inventory and routing decisions are taken sequentially is solved using the same model and method as in [1]. In the sequential scenario, various replenishment methods are employed for the inventory aspect, and a Large Neighbourhood Search (LNS) metaheuristic is utilized to optimize routing decisions from the UCC to the retailers. The integrated scenario is addressed using a matheuristic algorithm that combines mathematical optimization techniques with heuristics. A matheuristic algorithm based on the one presented in [2] is proposed to solve the integrated problem. We adapt the method to our specific problem context by extending it to a multi-product, heterogeneous, multi-trip context with a more complex objective function that also includes holding costs, order costs, and duration-based routing costs.

Computational results consistently demonstrate the superiority of the integrated approach across multiple performance metrics, including costs, logistics ratio, number of urban trips, loading degree, and distance traveled. A sensitivity analysis highlights critical factors, including retailer storage capacity, order cost, and retailer participation, that

influence the implementation of the integrated scenario. Retailers can benefit from the implementation of the integrated scenario to reduce their storage capacity requirements and allocate the available space for other purposes. Additionally, a higher order cost reduces the savings obtained from implementing the integrated scenario. Furthermore, higher levels of retailer participation result in more substantial cost reductions in both the second and first echelons.

Further analysis is conducted to investigate the impact of different transshipment costs on the cost difference between the sequential and integrated scenarios. As transshipment costs increase, the cost difference between the integrated and sequential scenarios becomes even larger.

References

[1] Iswari, T., Braekers, K., and Caris, A., (2023). Analyzing the benefits of a city hub: an inventory and routing perspective. *Computers & Industrial Engineering*, Vol 185, 109629.

[2] Bertazzi, L., Coelho, L.C., De Maio, A. and Laganà, D. (2019). A matheuristic algorithm for the multi-depot inventory routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 122:524–544.

The autonomy effect in warehousing – Insights from a large-scale lab experiment

T. De Lombaert, K. Braekers, K. Ramaekers
Hasselt University, Research Group Logistics
thomas.delombaert@uhasselt.be,
kris.braekers@uhasselt.be
katrien.ramaekers@uhasselt.be

R. De Koster
Erasmus University Rotterdam, Rotterdam School of Management
rkoster@rsm.nl

Keywords: Human factors ; Warehousing ; Job assignment

Warehouses play a prominent role in the supply chain of a company and contribute to its failure or success. Within a warehouse, many activities are performed, although it has been shown that order picking (OP) is by far the costliest. OP is defined as the process of retrieving products from their storage locations in a warehouse to prepare them for shipping. Since approximately 50% of the total warehousing costs can be attributed to OP, warehouse managers aim for high efficiency levels in the OP system. There are many planning decisions involved to devise an efficient OP system, for instance the batching of orders, the predetermination of picker routes, or the assignment of jobs/orders to pickers [1]. However, in many modern-day warehouses, a central planning system underpins these planning problems and sets out directives for human workers. This erodes the (perceived) autonomy of order pickers, although autonomy is one of the three basic psychological needs and found to affect worker well-being. The current study presents the testing and post-hoc evaluation of an autonomy-increasing experiment in a lab warehouse. We show that granting autonomy to order pickers has a beneficial impact on individual well-being as well as organisational outcomes.

One planning problem which contributes to efficient warehouse operations, is the job assignment planning problem. This planning problem manages the allocation of orders/batches/tasks (henceforward: orders) to pickers, as well as the sequence in which these orders should be picked. In many of today's warehouses, those assignments are completely random, apart from possible due dates (or priority) considerations. This results in the underutilisation of pickers' skills, as picker-order misfits can occur. For exam-

ple, a relatively short order picker could be assigned to top-shelf picks, or a picker with fear of heights might have to pick boxes at great height. Such mismatches can be avoided by repeatedly assigning pickers to specific orders, which are in line with their respective skills. However, previous research has highlighted the aversion of pickers from being assigned to products with the same characteristics over and over again [2]. In addition, this form of repetition may physically overstrain pickers, and lead to boredom as a result of monotony. The question arises whether it is possible to configure a system in which picker skills can be exploited and thereby also preserving the pickers' physical and mental well-being.

In this study, we investigate the impact of an order assignment mechanism in which order pickers get the opportunity to choose their next order from an order set that is presented to them at the depot. We expect that pickers will chose orders that are in line with their skills and/or preferences, resulting in a positive worker attitude. This assignment mechanism is tested during a three-month-long experiment. More precisely, we conduct a lab experiment in which we explore the impact of autonomous worker decisions. We also investigate the additional influence of incentive systems, hence a within-subjects 2x2 study design. In total, 165 students with an education background in logistics participate in our study. Results show that participants appreciate the possibility to shape their own work. Productivity, on the other hand, remains steady when comparing the system with and the one without worker autonomy. In addition, no strong moderation effects are found for the incentive system.

This study shows how high efficiency levels can coincide with increased worker involvement. Using a holistic evaluation approach, we show the beneficial impacts of an autonomy-increasing intervention. The insights derived from this study can be translated to other warehouse planning problems and give rise to several subproblems which ought to be encountered with an OR perspective.

References

[1] Van Gils, T., Ramaekers, K., Caris, A., & de Koster, R. B. (2018). Designing efficient order picking systems by combining planning problems: State-of-the-art classification and review. European Journal of Operational Research, 267(1), 1-15.

[2] De Lombaert, T., Braekers, K., De Koster, R., & Ramaekers, K. (2023). In pursuit of humanised order picking planning: methodological review, literature classification and input from practice. International Journal of Production Research, 61(10), 3300-3330.

Deep Reinforcement Learning for the Real-Time Inventory Rack Storage Allocation and Replenishment Problem

S. Teck
KU Leuven, KU Leuven Institute for Mobility - CIB
Sander.teck@kuleuven.be

T. S. Pham
Polytechnique Montréal, Department of Mathematics and Industrial Engineering

L.-M. Roussea
Polytechnique Montréal, Department of Mathematics and Industrial Engineering

P. Vansteenwegen
KU Leuven, KU Leuven Institute for Mobility - CIB

Keywords: Optimization, Storage Allocation and Replenishment, Logistics

The e-commerce sector is experiencing a rapid evolution towards increased automation [1]. The growing complexity of warehouse operations necessitates control systems capable of efficiently managing this complexity. Real-time scheduling, particularly in the context of inventory rack storage location and replenishment management, plays a crucial role in ensuring optimal performance. This study focuses on Robotic Mobile Fulfillment Systems (RMFS), a semi-automated warehousing system. Employing autonomous mobile robots (AMRs), the system eliminates the need for human operators within the storage area, with robots responsible for both storing and retrieving movable inventory racks. Human operators stationed at workstations alongside the storage area pick goods from these racks to fulfill customer orders. An intriguing characteristic of RMFS is its ability to dynamically alter the positioning of the inventory racks [2]. For instance, positioning may be based on the frequency of use of the inventory racks, allowing more frequently requested racks to be stored in locations closer to the workstations. This decision is important as it can significantly impact the overall system performance. Furthermore, during the warehouse activities, the stock within the racks will deplete. Thus, beyond deciding where to store inventory racks, determining the opportune time to restock the rack is essential, since a poorly designed system may accumulate large waiting queues

during operations.

The primary optimization objective considered in this study is to minimize the average cycle time of the mobile robots for both retrieval and replenishment activities. To address the real-time challenges of inventory rack storage allocation and replenishment in an RMFS, we propose a deep-reinforcement learning (DRL) approach. The decision-making agent in this study must make decisions following specific events, particularly after the fulfillment of retrieval requests. The storage area is partitioned into distinct storage zones due to the inherent challenges associated with explicitly allocating a storage location in a large warehouse layout [3]. The explicit allocation of storage locations would impede the learning process of the decision maker as it significantly increases its action space. By dividing the storage area into manageable zones, the agent is better equipped to explore and exploit the search space. The learning agent interacts with the environment, observes changes, and learns through trial-and-error from these interactions. In doing so, it is able to construct a policy by mapping environment states to actions. The action space include directly storing the inventory rack in a suitable storage zone or first replenishing the rack and subsequently storing it in a suitable storage zone. The proposed approach is extensively tested on a diverse set of problem instances varying in the number of storage locations, storage area layouts, and different storage zoning strategies. The efficiency of the intelligent decision maker is compared to the commonly used decision rules employed in the academic literature on this problem [4-6]. The experimental findings underscore the benefits associated with the integration of learning agent over the application of conventional decision rules in this warehouse environment, resulting in significant cycle time improvements. This research contributes to advancing the understanding of intelligent storage allocation and replenishment strategies for the real-time decision making processing within a semi-automated warehousing system.

References

[1] Boysen, N., de Koster, R., Weidinger, F. (2019). Warehousing in the e-commerce era: A survey. European Journal of Operational Research, 277, 396-411.

[2] Enright, J., Wurman, P.R. (2011). Optimization and Coordinated Autonomy in Mobile Fulfillment Systems. In Proceedings of the AAAI workshop on automated action planning for autonomous mobile robots, 33-38.

[3] Rimélé, A., Grangier, P., Gamache, M., Gendreau, M., Rousseau, LM (2021). E-commerce warehousing: learning a storage policy. arXiv: 2101.08828.

[4] Roy, D., Nigam, S., de Koster, R., Adan, I., Resing, J. (2019). Robot-storage zone assignment strategies in mobile fulfillment systems. Transportation Research Part E, 122: 119-142.

[5] Merschformann, M., Lamballais, T., de Koster, R., Suhl, L. (2019). Decision rules for robotic mobile fulfillment systems. Operations Research perspectives, 6: 100-128.

[6] Lamballais, T., Merschformann, M., Roy, D., de Koster, R., Azadeh, I., Suhl, L. (2022). Dynamic policies for resource reallocation in a robotic mobile fulfillment system with time-varying demand. European Journal of Operational Research, 300(3): 937-952.

Quantifying deviations in an order picking process through data-driven analysis

A. Leroy, K. Braekers, A. Caris
Hasselt University, Research Group Logistics
aicha.leroy@uhasselt.be
kris.braekers@braekers.be
an.caris@uhasselt.be

B. Depaire
Hasselt University, Research Group Business Informatics
benoit.depaire@uhasselt.be

T. van Gils
Atlas Copco, Airpower Service Center
teun.van.gils@atlascopco.com

Keywords: Order picking, Data-driven analysis, Warehousing

1 Abstract

Order picking is one of the most time-sensitive and cost-critical processes in warehousing, so researchers often focus on this warehouse operation. During this process, orders are retrieved from the warehouse storage to be shipped to customers. It is still applied manually in many warehouses, especially as the current trends, such as decreased product life cycle and increased product differentiation, amplify the need for flexible order picking systems and limit the possibility of automation. In these manual warehouses, picker-to-parts order picking is generally employed. Order pickers travel through the warehouse to collect requested products from the shelves following pre-determined routes. Researchers have developed numerous optimal and heuristic routing policies to guide order pickers efficiently. However, those routing policies often consider a single order picker in isolation or assume that multiple picker' routes are performed independently. In reality, numerous pickers travel through the warehouse simultaneously, and their routes can interfere. This phenomenon, called picker blocking, adds uncertainty to the picker routing problem and can lead to higher picking times than expected.

Using humans (i.e., order pickers) and their specific cognitive and physical skills to fulfill order picking activities makes the process more flexible. For example, when experiencing picker blocking, an order picker can independently decide to wait until the blockage is cleared or deviate from its planned route to avoid the blockage. On the other hand, these

human workers' decisions introduce uncertainties in the order picking process. Our goal is to learn about these uncertainties during the order picking process. These insights can be used to include more real-life factors in future research. Numerous aspects can influence the order in which order pickers complete a pick tour and the time it takes to complete the pick tour. We analyze deviations in the order picking process in a broad sense by focusing on both deviations in picking order and in time.

Related work on deviations during the order picking process is limited. First, Glock et al. (2017) use a qualitative approach by conducting surveys to find the types, causes, and consequences. The authors conclude that, although positive effects are possible, the consequences of deviations on order picking efficiency are mainly negative. Second, Elbert et al. (2017) compared optimal and heuristic routing strategies while considering route deviations. To do this, the authors use agent-based simulation to analyze the effects of route deviations on picking time under different routing methods. A range of artificially selected probabilities determines whether a route deviation occurs. They conclude that it is essential to consider route deviations when determining the preferred routing strategy.

We improve warehousing research by analyzing historical data to quantify the prevalence of order picking deviations. Our contribution to order picking research is twofold. Firstly, a methodology to extract insights about order picking deviations from a data set is proposed. Until now, the existence of order picking deviations has only been indicated in survey results. However, a method to extract these insights from real-life data sets is still missing. Secondly, the proposed methodology is applied to a case study to analyze the real-life impact of order picking deviations. Currently, no data-based insights about order picking deviations exist. Implementing real-life factors in order picking research is necessary to make sure that outputs from order picking planning models resemble reality more closely, and it also increases the likelihood of warehouse managers implementing scientific research on warehousing. As the effects of order picking deviations are expected to be mainly negative, the overall warehouse performance will no longer be overestimated. This results in expectations from warehouse managers that are attainable for order pickers and, therefore, higher employee well-being and lower burn-out rates may be achieved.

As order picker deviations may manifest in different ways, the aim is to identify both deviations from the planned pick order of items (e.g., locations in a pick aisle were skipped because the aisle was congested) and deviations from the expected times at which picks are performed (e.g., delays due to picker blocking or pickers taking alternative travel paths). Insights based on an extensive real-life data set will be presented.

2 References

[1] Elbert, R. M., Franzke, T., Glock, C. H., and Grosse, E. H. (2017). The effects of human behavior on the efficiency of routing policies in order picking: The case of route deviations. Computers & Industrial Engineering, 111, 537–551.

[2] Glock, C. H., Grosse, E. H., Elbert, R. M., & Franzke, T. (2017). Maverick picking: the impact of modifications in work schedules on manual order picking processes. International Journal of Production Research, 55(21), 6344-6360.

Determining well-being during a crisis based on social media data using natural language processing

X. Wang

Ghent University, Department of Marketing, Innovation, and Organization
Xiao.Wang@UGent.be

B. Janssens

Ghent University, Department of Marketing, Innovation, and Organization
FlandersMake@UGent–corelab CVAMO, Belgium
Research Foundation Flanders, Belgium
Bram.Janssens@UGent.be

M. Bogaert

Ghent University, Department of Marketing, Innovation, and Organization
FlandersMake@UGent–corelab CVAMO, Belgium
Matthias.Bogaert@UGent.be

L. Vanderbauwhede, L. Schetgen

Ghent University, Department of Marketing, Innovation, and Organization
lieselvanderbauwhede@hotmail.com, Lisa.Schetgen@UGent.be

Keywords: Social media analytics, Subjective well-being, NLP

The objective of this paper is to determine how subjective well-being (SWB) can be predicted based on Twitter (or X) data during a crisis. Having an accurate estimate of SWB based on social media can be important for government institutions to see how policies or other actions affect the population, without having to conduct time-consuming and expensive questionnaires. This is especially valuable during a crisis when rapid government responses are often crucial. Typically, a questionnaire, namely the International Positive and Negative Affect Schedule Short-form (I-PANAS-SF), is employed to assess SWB. As natural language processing (NLP) continues to advance, various techniques are employed to harness text data for predictive purposes [1].
Our study focuses on the COVID-19 pandemic as a representative crisis scenario. First, we obtain SWB by two methods: I-PANAS-SF (completed by human annotators) and sentiment analysis tools (i.e., valence aware dictionary for sEntiment reasoner (VADER) and TextBlob). The emotions associated with positive affect and negative affect, extracted from I-PANAS-SF, are individually labeled on a Likert scale. Second, we obtain structured features and unstructured features from our tweets. We evaluate five textual representation techniques (i.e., TF-IDF, GloVe, ELMo, RoBERTa, and OpenAI's ada embed-

ding) across five machine learning algorithms (i.e., elastic net regression, support vector machines, random forest, artificial neural networks, and eXtreme gradient boosting (XGBoost)) [2, 3]. Finally, we use metrics (i.e., $RMSE$ and R^2) to evaluate all the machine learning algorithms and identify the best combination of textual representation techniques and machine learning algorithms. We report $RMSE$ and R^2 of sentiment analysis tools to make a comparison with all prediction models. This study also attempts to provide insights into the most significant predictors of SWB prediction by assessing feature importances.

Results reveal that sentiment analysis tools prove to be suboptimal approaches for estimating SWB, as evidenced by the combined Alpayadin F-test with various prediction models. In contrast, prediction models constructed using the combination of textual representation techniques and machine learning algorithms exhibit relatively strong performance. Overall our empirical results show that models using advanced embedding methods (RoBERTa and OpenAI's ada embedding) are promising text representation techniques for predicting SWB. It is plausible that OpenAI's ada embedding holds the potential for superior performance compared to RoBERTa, albeit at the cost of increased computational time. The models using RoBERTa in our study (pretrained on social media data) appear to be less dependent on machine learning algorithms, which confirms the powerful ability of domain embedding for specific tasks. In a case study, we demonstrate that our methods allow for a more accurate tracking of SWB.

This research contributes to the existing literature by exploring two fundamental aspects: the suitability of using sentiment scores as SWB, and the optimal combination of text representation techniques and machine learning algorithms in the field of predicting SWB during a crisis based on social media data.

References

[1] Janssens, B., Schetgen, L., Bogaert, M., Meire, M., & Van den Poel, D. (2023). 360 Degrees Rumor Detection: When Explanations Got Some Explaining To Do. European Journal of Operational Research.

[2] Borchert, P., Coussement, K., De Caigny, A., & De Weerdt, J. (2023). Extending business failure prediction models with textual website content using deep learning. European Journal of Operational Research, 306(1), 348-357.

[3] Stevenson, M., Mues, C., & Bravo, C. (2021). The value of text for small business default prediction: A deep learning approach. European Journal of Operational Research, 295(2), 758-771.

Decoding Stock Market Reactions to Earnings Call Dialogues

P. Borchert
IESEG School of Management, KU Leuven

K. Coussement
IESEG School of Management

J. De Weerdt
KU Leuven

A. De Caigny
IESEG School of Management

Keywords: Natural Language Processing, Finance, Predictive Analytics

1 Abstract

This research investigates the influence of earnings conference call Q&A sections on stock market responses. By harnessing the analytical capabilities of Language Models, we analyze earnings conference call transcripts from a broad spectrum of publicly traded companies. Our objective is to determine the extent to which these Q&A interactions predict abnormal stock returns. Our experiments demonstrate that the Q&A segments provide substantial predictive power. This provides insights for understanding the role of corporate communication in shaping market perceptions and influencing stock performance. The findings are particularly relevant for investors, analysts, and corporate communicators, highlighting the importance of the information exchanged during these interactive sessions.

2 Methodology and Preliminary Findings

Our approach utilizes Language Models to analyze transcripts of earnings conference calls from leading US stock exchanges, namely NASDAQ and NYSE. We place a special focus on the Q&A section, a segment distinguished by interactive dialogue format. In this part of the call, executives engage with and respond to inquiries from analysts and investors. Our comprehensive analysis indicates that the content of earnings disclosures, particularly the Q&A sections, provides predictive insights for market reactions. We note that the predictive value of different sections within these transcripts is not uniform across various industries. Our preliminary findings suggest a spectrum of information content in the Q&A segments. This ranges from repetition of details already covered

Tell Me a Story! Narrative-Driven XAI with Large Language Models

D. Martens
University of Antwerp

J. Hinns
University of Antwerp

Keywords: Explainable AI, Counterfactual Explanations, SHAP, Large Language Models, XAIStories.

The dominance of complex black-box machine learning models for most state-of-the-art prediction tasks, highlights the need for Explainable AI (XAI). Whilst feature attribution methods such as SHAP and LIME are frequently used to quantify feature importance, their numerical format can pose understanding challenges for lay-users. Additionally, while counterfactual (CF) explanations offer 'what if' scenarios, they often leave users struggling to understand the 'why' behind them. To mitigate these issues, we present XAIstories. Through the use of Large Language Models, XAIstories provide narratives of possible rationale behind AI predictions: We present two implementations – SHAPstories, based on SHAP explanations of a prediction score, and CFstories, based on CF explanations of a decision. Our findings present a compelling case: upwards of 90% of respondents from the general public find the narratives created by SHAPstories to be persuasive. Data scientists identify the primary value of SHAPstories in its ability to communicate explanations to a general audience, with 92% of data scientists indicating that it will contribute to the ease and confidence of lay-users in understanding AI predictions. 83% of data scientists indicate they are likely to use SHAPstories for this purpose. In image classification, CFstories are considered at least as convincing as users own crafted stories by over 75% of lay-user participants. Furthermore, CFstories improves accuracy by over 20% and is ten times faster than manually created narratives. Our results suggest that XAIstories empower non-data scientists to better grasp complex AI predictions. By providing explanations that appeal to a broader audience, we offer domain experts, especially those without knowledge of data science, a platform to evaluate models with greater ease and confidence based on their explanations.

Leveraging zero-shot learning for museum collection data enrichment

M. Reusens
KU Leuven, LIRIS
manon.reusens@kuleuven.be

A. Adams
National Museum of the Royal Navy

R. Pacheco-Ruiz
National Museum of the Royal Navy

S. vanden Broucke
UGent, Department of Business Informatics and Operations Management
KU Leuven, LIRIS

B. Baesens
KU Leuven, LIRIS
University of Southampton,
School of Management

Keywords: keyword extraction, keyword enrichment, chatbots

Nowadays, large language models have gained widespread popularity. The substantial performance increase in these models coincides with the accumulation of external knowledge. This can be leveraged for multiple purposes, facilitating the automation of tasks that were previously handled manually. In the domain of museum collection management, this automation allows individuals to focus on challenging tasks that require human expertise. More specifically, this is well suited for tasks such as keyword extraction and enrichment regarding descriptions of historical objects when extensive descriptions are available.

This research investigates the utility of leveraging a large language model as a zero-shot learner for keyword extraction and enrichment based on descriptions of historical objects. Employing a unique dataset from the National Museum of the Royal Navy, we investigate the model's capability to accurately extract keywords directly from descriptions of historical objects provided by curators. Additionally, we research whether the model can correctly apply external knowledge to enhance keyword extraction and autonomously augment the set of relevant keywords. Our results are manually evaluated by a team of experts allowing us to meticulously examine patterns within the results, particularly instances where external knowledge was (not) accurately employed or where too detailed results were obtained. This work contributes to the broader understanding of the

Improving the robustness of a railway timetable in a bottleneck area: a case study in Belgium

I. Van Hoeck
KU Leuven, Institute for Mobility - CIB
inneke.vanhoeck@kuleuven.be

P. Vansteenwegen
KU Leuven, Institute for Mobility - CIB
pieter.vansteenwegen@kuleuven.be

Keywords: Railway transport, Timetable robustness, Timetabling

The Belgian railway network is very dense and heavily used. This can easily cause delays to propagate through the network, which obviously has a negative effect on the travel time of the passengers. Therefore, we aim to improve the passenger robustness, which is defined as minimizing the total travel time of all passengers in practice in case of frequently occurring small delays (Dewilde et al, 2011). However, calculating this objective directly is computationally expensive. Instead, Dewilde et al. (2014) and Burggraeve and Vansteenwegen (2017) indirectly strive for passenger robustness by looking for a good spreading in time of the trains by considering the buffer times between them. Caimi et al. (2011) also argue that maximizing the buffer times between trains will allow small delays to be absorbed. Our objective function is based on their work and focuses on improving the smallest buffer times. Between each pair of trains that have infrastructure in common, the buffer time is defined as the minimum time span between the two trains. A piecewise linear, monotone decreasing function is used to associate a cost with a buffer time. The slope of these linear pieces becomes less steep as the buffer times become larger. This corresponds to the idea that it is much more valuable to increase a (very) small buffer time, than to increase a large buffer time even more. The objective is to minimize the sum of the costs for all train pairs.

The timing, the dwell times and the microscopic routing of the trains can be adapted to improve the robustness. Besides being conflict-free on a microscopic level, the timetable should satisfy certain properties to be useful in practice for the Belgian railway companies. Therefore, an LP model is developed where the user can decide which constraints to impose depending on the properties that the resulting timetable should have. The following properties can be included: fixing the timing of a train, providing a time window for the timing of a train, spreading in time between two trains, symmetry between two trains and the reuse of a train.

A case study of a bottleneck area centered around the station of Halle is considered. NMBS and Infrabel provided the necessary information about the timetable properties

that should hold. The most restrictive of these properties is related to the timing of the trains. Because the considered area is located just outside of the even larger bottleneck of Brussels, additional constraints are required to guarantee that the new timetable is also feasible in that part of the network. This is done by fixing the timing of the trains at the border with Brussels. Several experiments are performed for this case study. First, the timetable is optimized while all the necessary properties are imposed. An improvement of 18% in the objective value compared to the original timetable is found. Additional experiments with a varying number of alternative routing options and more freedom for the timing at the border with Brussels are performed as well. The results give more insight into which changes have the most impact on the quality of the timetable. Finally, the model is used to analyze and optimize alternative scenarios that are provided by NMBS and Infrabel. One of these scenarios looks at the addition of trains. The results show that two trains can actually be added to the timetable, while still obtaining a better robustness than for the current timetable. Another scenario considers a list of possible trains that can be cancelled. Obviously, cancelling a train will always lead to an improvement in the robustness, but the results clearly show that cancelling certain trains leads to a larger improvement than others.

References

[1] Burggraeve, S., Vansteenwegen, P., 2017. "Robust routing and timetabling in complex railway stations", Transportation Research Part B, vol. 101, pp. 228-244.

[2] Caimi, G., Chudak, F., Fuchsberger M., Laumanns, M., Zenklusen, R., 2011. "A new resource-constrained multicommodity flow model for conflict-free train routing and scheduling", Transportation Science, 45, pp. 212-227.

[3] Dewilde, T., Sels, P., Cattrysse, D., Vansteenwegen, P., 2011. "Defining robustness of a railway timetable", In Proceedings of the 4th international Seminar on Railway Operations Modelling and Analysis (RailRome2011), Rome, Italy.

[4] Dewilde, T., Sels, P., Cattrysse, D., Vansteenwegen, P., 2014. "Improving the robustness in railway station areas", European Journal of Operational Research, vol. 235, pp. 276-286.

Capacitated Examination Timetabling: Multi-Neighbourhood Simulated Annealing

D. Van Bulck, D. Goossens
Ghent University, Department of Business Informatics and Operations Management
david.vanbulck@ugent.be, dries.goossens@ugent.be

A. Schaerf
University of Udine, Polytechnic Department of Engineering and Architecture
andrea.schaerf@uniud.it

Keywords: Metaheuristics, Neighbourhoods, Simulated Annealing

1 Introduction and problem description

Many variations of the Examination Timetabling Problem exist, each with their own resources, constraints, and objectives. This talk focuses on the classic version proposed for the Second International Timetabling Competition (ITC-2007; see McCollum et al. (2010) and Ceschia et al. (2022)). Here, the task involves allocating each exam to a single period and room, allowing for room sharing among exams. Prominent constraints include ensuring each student has only one exam per period (i.e., the timetable is 'conflict-free') and that, for each period, the capacity of a room is sufficient to host all students of the exams assigned to it. The objective is to distribute student exams evenly over time, avoiding scheduling exams with different durations in the same room, and considering room and period preferences for specific exams.

2 Solution method

The literature on examination timetabling almost exclusively focusses on the design of metaheuristics, ranging from standard techniques like simulated annealing and tabu search, to more exotic variants like harmony search and intelligent water drops. Despite the plethora of methods proposed, we observe that almost all neighbourhoods used for ITC-2007 originate from uncapacitated timetabling and have been straightforwardly adapted to deal with the assignment of exams to rooms. Moreover, we observe that neighbourhoods are often restricted to the most basic ones as it is not always clear how to (efficiently) implement more complex ones like Kempe chains. In addition to presenting an overview of existing operators and their implementation, we introduce two novel

neighbourhoods coined $Beam$ and $Component$. Given a subset of exams assigned to a single period, $Beam$ re-optimizes the room assignment via beam search. On the other hand, $Component$ exploits the fact that most conflict graphs associated with ITC-2007 instances turn out to be disconnected. Hence, stretching the interval during which exams of a disconnected component are scheduled may enhance the spread for students, without introducing any new student conflicts.

We integrate five existing as well as the two newly developed neighbourhoods into the framework of multi-neighbourhood Simulated Annealing (SA). Our search space is composed by an array of pairs that assigns to each exam a period and a room, and also includes solutions that may violate hard constraints such as conflicts or room capacities. These violations are included in the cost function, along with the soft constraints, but with a suitably larger weight. During the development of the algorithm, we avoided premature commitment to certain design choices by including these choices as parameters of the algorithm (a paradigm known as programming by optimization). All parameters of the algorithm were then tuned with `irace`, followed by an extensive ablation analysis to assess the contribution of the various neighbourhoods.

3 Experimental results

To assess the efficacy of our algorithm, we executed it twice: once with a short runtime (less than five minutes, adhering to the official ITC-2007 time limit), and once with an extended runtime of approximately 7 hours (100 times the official ITC-2007 time limit, which is more in line with Gogos et al. (2010)). With regard to the official time limit, our algorithm turns out to be competitive to or even outperforms existing state-of-the-art, even if restricted to the use of the most basic yet efficiently-implemented neighbourhood operators. Moreover, using the extended runtime, our algorithm finds new best solutions for about half of the problem instances. These instances have been validated and uploaded to OptHub (see `https://opthub.uniud.it/problem/timetabling/edutt/ett/itc-2007-ett`). For a comprehensive analysis of our results, please refer to our published work in Van Bulck et al. (2023).

References

[1] Ceschia, S., Di Gaspero, L., & Schaerf, A. (2022). Educational timetabling: Problems, benchmarks, and state-of-the-art results. European Journal of Operational Research, 308, 1-18.

[2] Gogos, C., Goulas, G., Alefragis, P., Kolonias, V., & Housos, E. (2010). Distributed scatter search for the examination timetabling problem. In B. McCollum, E. K. Burke, & G. White (Eds.), 8th International Conference on the Practice and Theory of Automated Timetabling (PATAT-2010) (pp. 211–223).

[3] McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A. J., Di Gaspero, L., Qu, R., & Burke, E. K. (2010). Setting the Research Agenda in Automated Timetabling: The Second International Timetabling Competition. Informs Journal on Computing, 22, 120-130.

[4] Van Bulck, D., Goossens, D., & Schaerf, A. (2023). Multi-neighbourhood simulated annealing for the ITC-2007 capacitated examination timetabling problem. Journal of Scheduling, In Press.

Reforming the Belgian youth field Hockey Leagues with an incomplete round-robin tournament

K. Devriesere
Ghent University, Department of Business Informatics and Operations Management
karel.devriesere@ugent.be

D. Goossens
Ghent University, Department of Business Informatics and Operations Management,
Flanders Make, CVAMO Core Lab
dries.goossens@ugent.be

Keywords: Sports timetabling, incomplete round-robin, variable neighborhood search

Traditionally, teams from the youth leagues of the Royal Belgian Hockey Association (RBHA) are partitioned into groups. Each group is then organized as a double round robin tournament, meaning that each team meets each other team in its group exactly twice. The groups are formed based on the travel distance between the teams. Given the assignment of teams to groups, a schedule is constructed for each group.

In this traditional competition design, the opponent set and travel distance of each team are determined through the assignment of teams to groups. We propose to organize each of the U7-U9 hockey leagues according to a more flexible design, namely an incomplete round robin tournament (Cochran, 1971). Instead of partitioning the teams into different groups, in this new format each team's opponent set is the whole league. A league is understood as a set of teams that share the same age, gender and strength level. Strength levels are determined by the RBHA based on the input of the clubs. Inter-league matches are forbidden. We further restrict teams from facing the same opponent more than twice. There are interdependent capacity constraints, since hockey fields are shared among teams in different leagues. Since the incomplete round robin design offers more flexibility with respect to the possible opponent sets of each team, this competition design has great potential to improve the total travel distance of the leagues.

The goal of professional sports competitions is typically to reveal the relative strength of teams (Placek, 2023). For this reason, an integral part of professional sports competitions is the ranking method, which produces an ordering of the teams based on pairwise comparisons (Vaziri et al., 2017). Because each team potentially faces very different opponents, constructing a fair ranking method for an incomplete instead of a traditional round robin tournament is considerably more difficult. In youth sports leagues, however, revealing the true strength of teams is only of secondary importance. Instead, other fac-

tors related to development, social interaction and enjoyment of sports are deemed far more important in youth leagues (Wankel & Kreisel, 1985). In fact, for age categories U7-U9, the RBHA even desires to omit an official ranking so as to fully focus on the development of the players. Hence, while developing a fair ranking is considerably more complicated for the incomplete round robin design, we can safely ignore this concern for the hockey youth leagues.

We propose an integer program formulation to solve this problem. Since the instance size (> 800 teams) makes it impractical to find good solutions in a reasonable amount of time by directly solving the model with a commercially available software, we developed a Variable Neighborhood Search to tackle this problem. A neighborhood is generated by fixing a subset of the variables while releasing all other variables. This heuristic is known as Fix-and-Optimize and has recently proven successful in a number of sports scheduling applications (Van Bulck & Goossens, 2020). We show that the incomplete round robin design performs much better with respect to travel time compared to the traditional design used by the RBHA. Furthermore, we show that our heuristic is able to give solutions close to optimality within a reasonable amount of time.

References

[1] Cochran, R. S. (1971). Designs with redeeming social aspects for evenings of social bridge. The American Statistician, 25(2), 12-15.

[2] Placek, P. C. (2023). The impossibility of a perfect tournament. Entertainment Computing, 45, 100540.

[3] Van Bulck, David; Goossens, Dries. (2020). The international timetabling competition on sports timetabling (ITC2021). European Journal of Operational Research, 2023, 308.3: 1249-1267.

[4] Vaziri, B., Dabadghao, S., Yih, Y., & Morin, T. L. (2018). Properties of sports ranking methods. Journal of the Operational Research Society, 69(5), 776-787.

[5] Wankel, L. M., & Kreisel, P. S. (1985). Factors underlying enjoyment of youth sports: Sport and age group comparisons. Journal of Sport and Exercise Psychology, 7(1), 51-64.

Multi-league sports scheduling

D. Goossens
Ghent University, Department of Business Informatics and Operations Management
Flanders Make, CVAMO Core Lab
dries.goossens@ugent.be

M. Li
Ghent University, Department of Business Informatics and Operations Management
miao.li@ugent.be

M. Davari
SKEMA Business School, Centre for Analytics and Management Science
morteza.davari@skema.edu

Keywords: sports timetabling, venue capacity, travel distance

1 Introduction

Sports scheduling, i.e. deciding which match is to be played when, is an important part of any sports competition. Various papers describe the scheduling of a single league as a highly challenging problem, which requires reconciling various wishes constraints from several stakeholders using advanced methods (see Van Bulck et al. (2020) and Ribeiro et al. (2023) for an overview). In amateur and youth sports, however, the sports scheduling problem typically involves multiple leagues, which are categorized based on the age or skill of the players. This number of leagues may be several hundreds, involving thousands of teams, and tens of thousands of matches to be scheduled. Moreover, since teams from the same club share the same venue, a capacity problem at each club results. Indeed, if a team is the home team for its match, its club should host the match, and the maximum number of matches a club can host simultaneously is limited by its number of terrains. We say that a club has a venue capacity violation if at some point in time, it has more teams that are scheduled a home game than it can accommodate. As the teams of a club play in different leagues, this capacity problem creates dependencies between the leagues.

2 A step-wise approach

To handle multi-league tournaments in practice, league organizers typically use a step-wise approach. First, the teams are grouped into leagues. The main goal here is to mini-

mize the resulting total travel distance while respecting competitive balance constraints. This problem is known as the sports team grouping problem (see e.g. Toffolo et al. (2019). In order to avoid collusion and match-fixing, league organizers strive to compose leagues such that they have at most one team per club. Second, the organizer decides on a starting round for each league, and next determines for each team on which round it plays a home game. This is reasonable, as the club's venue capacity (and possibly also availability) is the most important constraint in youth and amateur sports. It is done by assigning each team to a so-called Home-Away Pattern (HAP), which indicates for each round whether its team has a home game (H), or an away game (A). Finally, combining the HAP assignment with a compatible opponent table, which specifies each team's opponent for each round, the schedule follows. Note that not just any set of HAPs will allow a schedule. Despite being a well-researched topic (see e.g. Briskorn (2008), Van Bulck and Goossens (2020)), the complexity of deciding whether a schedule exists for a given HAP set is still open.

3 Results

In this talk, we present a heuristic to handle large-scale multi-league scheduling problems. We illustrate the impact of using different sets of home-away patterns on the total venue capacity violations. We also discuss the trade-off between minimizing the total distance travelled by all teams and minimizing the total venue capacity violations over all clubs.

4 References

[1] Briskorn, D. (2008), Feasibility of home away pattern sets for round robin tournaments, Operations Research Letters, 36(3), 283–284.

[2] Ribeiro, C. C., Urrutia, S., and de Werra, D. (2023). A tutorial on graph models for scheduling round-robin sports tournaments. International Transactions in Operational Research, 30(6):3267–3295.

[3] Toffolo, T. A., Christiaens, J., Spieksma, F. C. and Berghe, G. V. (2019), The sport teams grouping problem, Annals of Operations Research 275(1), 223–243.

[4] Van Bulck, D. and Goossens, D. (2020), On the complexity of pattern feasibility problems in time-relaxed sports timetabling, Operations Research Letters 48(4), 452–459.

[5] Van Bulck, D., Goossens, D., Schönberger, J. and Guajardo, M. (2020), RobinX: a three-field classiffication and unified data format for round-robin sports timetabling, European Journal of Operational Research 280(2), 568–580.

ManyWorlds: Combinatorial Programming with Functions

J. Devriendt
Nonfiction Software
jo.devriendt@nonfictionsoftware.com

Keywords: Combinatorial programming, integer programming, AMPL.

In its most abstract form, *combinatorial programming* consists of writing down unambiguous expressions that a computer can interpret and solve. Those expressions range from simple propositional clauses (SAT solving) over mathematical (in)equalities (pseudo-Boolean solving, mixed integer programming (MIP)...) to algebraic or logic-based formalisms (mathematical programming, first order logic and extensions, answer set programming, satisfiability modulo theories...), perhaps with special purpose constraints (constraint programming (CP)).

The core of all these formalisms, which this paper unifies under the notion of combinatorial programming, is that a set of *possible worlds* (solutions) is described by an unambiguous expression, and that the backend solvers, grounders, compilers and translators for these formalisms, reason about this set of possible worlds. The focus of a formalism or a backend may differ: e.g., it may specialize in finding an optimal world (optimization), on deciding whether at least one world exists (satisfiability), on counting the number of worlds, on finding the intersection of all worlds (propagation), on explaining logical consequences etc. But in the end, this all comes down to reasoning about a set of possible worlds.

ManyWorlds is a new high-level combinatorial programming language that centers its syntax and semantics squarely on the concept of a function, which both programmers and high-schoolers are familiar with. Additionally, in an effort to be accessible, ManyWorlds styles itself as a classic programming language akin to Python, with simple keywords for builtin operators and functions, using short-hand symbols only for basic well-known arithmetic. It also introduces the *fold-map-filter* expression, which, similar to set and list comprehensions, greatly simplifies specifying aggregate expressions.

Beyond its syntax and semantics, ManyWorlds aims to provide a beginner-friendly user experience. A no-install online editor is available on manyworlds.site with basic syntax highlighting and in-browser syntax checking. An effort is made to provide clear syntax errors and advanced debug information in case of unsatisfiability. ManyWorlds is free and open with source code available at gitlab.com/nonfiction-software/manyworlds.

ManyWorlds' main inspiration is FO(·) [1]: an extension of typed first order logic with arithmetic, aggregates and inductive definitions. ManyWorlds aims to be a simpler, less mathematical and more approachable language than FO(·), using functions as basic building block instead of FO(·)'s relations. OPL [6], Essence [4] and MiniZinc [5] are high-level constraint programming languages similar to ManyWorlds. The main difference is the simplicity and accessibility ambition of ManyWorlds – i.a., a user does not need to learn about variable arrays, matrices, global constraints, decision variables, parameters…

The same argument applies to high-level mathematical programming languages such as AMPL [3]. For a particular example where AMPL requires a user to do more mental work than needed, consider a `next` function mapping a city to the next in an unknown Hamiltonian cycle (e.g., for a TSP). The following ManyWorlds expression elegantly captures the total distance objective:

```
sum [ distance(x,next(x)) for x where City(x) ]
```

The ManyWorlds compiler removes the burden to a user of unnesting this `next` function and automatically translates this to the usual weighted linear sum:

$$\sum_{x \in City, y \in City} Distance(x, y) * (next(x) = y)$$

In AMPL syntax, one could conceivably specify

```
sum {x in City} Distance[x,next[x]]
```

but this ultimately fails as "[v]ariables in subscripts are not yet allowed".

Similarly to AMPL and other high level combinatorial languages, ManyWorlds translates the high level specification to a low level integer program solved by its backend solver Exact [2]. Performance-wise, ManyWorld's goal is not to beat state-of-the-art MIP or CP solvers, but rather, to provide decent performance on a wide range of computations, including the less well studied intersection and explanation computations.

References

[1] Cat, B.D., Bogaerts, B., Bruynooghe, M., Janssens, G., Denecker, M.: Predicate logic as a modeling language: the IDP system. In: Kifer, M., Liu, Y.A. (eds.) Declarative Logic Programming: Theory, Systems, and Applications, pp. 279–323. ACM / Morgan & Claypool (2018). https://doi.org/10.1145/3191315.3191321

[2] Devriendt, J.: Exact solver (2023), https://gitlab.com/JoD/exact

[3] Fourer, R., Gay, D.M., Kernighan, B.W.: A modeling language for mathematical programming. Management Science **36**(5), 519–554 (1990), http://www.jstor.org/stable/2632268

[4] Frisch, A.M., Harvey, W., Jefferson, C., Hernández, B.M., Miguel, I.: Essence: A constraint language for specifying combinatorial problems. Constraints **13**, 268–306 (2007), https://api.semanticscholar.org/CorpusID:1931705

[5] Nethercote, N., Stuckey, P.J., Becket, R., Brand, S., Duck, G.J., Tack, G.: MiniZinc: Towards a standard CP modelling language. In: Bessière, C. (ed.) Principles and Practice of Constraint Programming – CP 2007. pp. 529–543. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)

[6] Van Hentenryck, P.: The OPL Optimization Programming Language. MIT Press, Cambridge, MA, USA (1999)

Parametric upper and lower bounds of linear variations of a linear problem's LHS

B. Miftari
University of Liege,
Montefiore Institute
bmiftari@uliege.be

G. Derval
University of Liege,
Montefiore Institute
gderval@uliege.be

Q. Louveaux
University of Liege,
Montefiore Institute
q.louveaux@uliege.be

D. Ernst
University of Liege,
Montefiore Institute
dernst@uliege.be

Keywords: Parametric linear optimization, bounds, left-hand side.

1 Abstract

For many linear problems, a lot of assumptions are made upon uncertain constraint coefficients. These uncertainties can be lifted by performing an in-depth analysis of the optimum. One of the tasks related to this analysis consists of assessing the impact of the assumptions taken on the final objective function value. In this presentation, we focus on linear modification of the constraint matrix. For a given parameter λ, the problem can be written as

$$\begin{aligned} f(\lambda) = \min \quad & c^t x \\ \text{s.t} \quad & A_1 x \leq b_1 \\ & A_2 x + \lambda D x \leq b_2. \end{aligned} \tag{1}$$

The matrix D contains all the coefficients impacted by the parameter λ. In the literature, a lot of work has been done regarding modifications on the right-hand side [1, 5, 2] and objective [7, 4, 9] but little on the left-hand side. Lifting the uncertainty on the left-hand side relies on heavy computations, re-optimizations or approximations to compute the objective for every value of λ [3, 8, 6]. In order to alleviate this issue, in this paper, we propose a novel approach consisting of finding upper and lower bounds of the problem depicted in (1) that require fewer computations and give guarantees upon the behavior of the objective. We formulate several bounds based on three techniques:

- Optimum reuse: we optimize the problem and reuse the optimal as long as it stays feasible.
- Robust optimization: We reformulate the inner problem using several robust optimization approaches separated in two categories: one consisting of constant solution for a given range, the others as an affine function of λ (resulting, respectively, in a constant and a linear bound). These techniques applied on the primal problem provide upper bounds.
- Lagrangian relaxation: we relax the constraints linked to the modification, $A_2x + \lambda Dx \leq b_2$, while adding the term $A_2x + \lambda Dx - b_2$ to the objective function c^tx with an associated cost. This relaxation on the primal provides lower bounds. From this principle, we derive three bounds of increasing complexity: a constant one, a linear one, and a quadratic one.

Their application to the dual problem leads to the opposite type of bound. The bounds provide guarantees in between the computed points, i.e. they do not allow for any outliers, and require a smaller number of computations. They give an indication of the evolution of the objective and can help practitioners to focus on particular values of λ within the interval that have interesting or unconventional behavior. Moreover, a basic iterative refinement algorithm allows computing the function $f(\lambda)$ up to a given precision using the bounds. We illustrate the bounds on two problems from the energy field and provide benchmarks.

References

[1] Ilan Adler and Renato D. C. Monteiro. A geometric view of parametric linear programming. *Algorithmica*, 8:161–176, 1992.

[2] Thomas Gal and Josef Nedoma. Multiparametric linear programming. *Management Science*, 7(18):406–422, 1972.

[3] Tomas Gal. *Postoptimal Analyses, Parametric Programming, and Related Topics, Degeneracy, Multicriteria Decision Making, Redundancy.* De Gruyter, Berlin, New York, 1994.

[4] S. I. Gass and Thomas L. Saaty. Parametric objective function (part 2)- generalization. *Journal of the Operations Research Society of America*, 3(4):395–401, 1955.

[5] A. M. Geoffrion and R. Nauss. Parametric and postoptimality analysis in integer linear programming. *Management Science*, 23(5):453–466, 1977.

[6] Rajab Khalilpour and I.A. Karimi. Parametric optimization with uncertainty on the left hand side of linear programs. *Computers & Chemical Engineering*, 60:31–40, 2014.

[7] Thomas Saaty and Saul Gass. Parametric objective function (part 1). *Journal of the Operations Research Society of America*, 2(3):316–319, 1954.

[8] Jack Sherman and Winifred J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.

[9] Rob A. Zuidwijk. Linear parametric sensitivity analysis of the constraint coefficient matrix in linear programs. *ERIM REPORT SERIES RESEARCH IN MANAGEMENT*, ERS-2005-055-LIS, 2005.

Residual Simplex method with applications in Network Design

W. Vanroose
Universiteit Antwerpen, Department of Mathematics
wim.vanroose@uantwerpen.be

Keywords: Network flow, Dual Simplex, subspace methods

1 Subspace methods

Krylov subspace methods such as *Conjugate Gradients* (CG) and *Generalized Minimal residual* (GMRES) solve extremely large linear systems in parallel. They are used as solvers in many engineering and natural sciences simulations for problems such as fluid flow, quantum mechanics and mechanical structure calculations [1]. These methods work by creating a subspace of vectors spanned by the residuals of the subsequent guesses for the solution. These residual vectors form an orthogonal basis. This leads to fast convergence since these methods never search in the same direction. They can solve problems with 100s of millions of variables and avoids the factorization of the sparse matrices that appear in these problems.

2 Residual Simplex method

Inspired by these subspace methods and ResQPASS [2] we propose the *Residual Simplex* method, where a sparse linear programming problem (LP) of the form $\min c^T x$ subject to $Ax = b$ and $l \le x \le u$ is projected on the basis of residual vectors, where $c, x, l, u \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m\times n}$. The iterations are defined as

Definition 2.1 *The residual subspace LP iteration is a sequence $x_k \in \mathbb{R}^n$ of guesses with $k = 1, 2, \ldots$ where*

$$\begin{aligned} x_k = \operatorname{argmin} & \, c^T x \\ \text{s.t. } & V_k^T (Ax - b) = 0 \\ & l \le x \le u, \end{aligned} \tag{1}$$

where

$$V_k = \operatorname{span}\left\{ \frac{r_0}{\|r_0\|}, \frac{r_1}{\|r_1\|}, \ldots, \frac{r_{k-1}}{\|r_{k-1}\|} \right\}, \tag{2}$$

where $r_l = Ax_l - b \in \mathbb{R}^m$ and x_0 is any initial feasible solution for the constraints.

The resulting method is a combination of constraint aggregation and a subspace method. We solve the sequence of problems with a dual simplex [3]. Each new problem can be warm-started with the previous solution. The method is especially performant for strongly degenerate problems, where the basic solution hits the lower and upper bounds.

In this presentation, we introduce the method, derive its most important properties and illustrate the performance for network design problems from the airline industry.

3 References

[1] Liesen, Jörg, and Zdenek Strakos. *Krylov subspace methods: principles and analysis.* Numerical Mathematics and Sciences, 2013.

[2] Symoens, Bas, and Wim Vanroose. *Residual QPAS subspace (ResQPASS) algorithm for bounded-variable least squares (BVLS) with superlinear Krylov convergence.* arXiv preprint arXiv:2302.13616 (2023).

[3] Koberstein, Achim. *Progress in the Dual Simplex algorithm for solving large scale LP problems: techniques for a fast and stable implementation* Computational Optimization and Applications 41 (2008): 185-204.

Grid-based tow train routing optimization for efficient part feeding

G. Gündüz Mengübaş
University of Antwerp, Department of Engineering Management, ANT/OR
gul.gunduzmengubas@uantwerpen.be

K. Sörensen
University of Antwerp, Department of Engineering Management, ANT/OR
kenneth.sorensen@uantwerpen.be

Keywords: narrow aisles, optimization, routing

1 Abstract

Effective part delivery to an assembly line plays a crucial role in Just-in-Time (JIT) manufacturing, and is frequently accomplished using tow trains. Nonetheless, the optimization of tow train routing becomes extremely important, especially when dealing with the complexities of intricate and adaptable manufacturing systems.

This research focuses on the need to attain unobstructed (blocking free) tow train routing in manufacturing settings with limited aisle space. The suggested grid-based approach for tow train routing utilizes optimization techniques that take into account blocking constraints.

This research makes several contributions. First, we divide the layout of the production facility into grid-based pixels, and implement Dijkstra's algorithm to find the shortest path between workstations. The pixel layout allows us to easily verify whether tow trains block each other, i.e., occupy the same pixel at the same time. Second, we model the problem of delivering parts to workstations as a Generalized Vehicle Routing Problem (GVRP). Third, we develop a Simulated Annealing (SA) algorithm to solve the GVRP. Fourth, we implement several operators to avoid tow train blocking by rerouting the tow trains when conflicts are detected on the grid-based layout.

Mixed-model assembly sequencing with workers' competencies

C. Miguel, S. Hoedt, J. Cottyn, E.H. Aghezzaf
Ghent University, Department of Industrial Systems Engineering and Production Design
Flanders Make@UGent, Industrial Systems Engineering
carlos.miguelpalao@ugent.be; steven.hoedt@ugent.be
johannes.cottyn@ugent.be; elhoussaine.aghezzaf@ugent.be

Keywords: Mixed-Model Assembly, Competence, Sequencing

1 Introduction

Currently, customers have access to a wide portfolio of variants of the same product, called models. These are tailored towards different customers' desires, thus sharing a main frame of characteristics while differing in others. These models are normally manufactured in mixed-model assembly lines (MMAL). In any given workstation, the total processing time changes with the sequence of products through the line. This can disrupt the flow of the line if several high processing time models are scheduled consecutively [1], leading to work that can not be completed in its assigned slot, called work overloads. This type of assembly line features a wider variety of tasks than their conventional counterparts. Manual assembly is thus more common, given the higher adaptability character of manual workers compared to automated systems. Despite this, workers may have differing competence levels, which may result in varying performance [2] depending on the tasks they are assigned to.

Worker assignment and product sequencing in manual assembly are treated in literature as separate problems. Nevertheless, these two decisions are linked. Unforeseen inefficiencies may occur if operators are assigned to tasks for which they lack the required competence. Moreover, assignation priorities may also change depending on the product mix or on the competence of the workforce. Our research presents a proposal to extend the current models to consider both problems simultaneously.

2 Problem extension

We start with the mixed-model sequencing problem (MMSP) proposed by Scholl [3], specifically in its formulation minimizing work overloads, MMSP-W. This problem is formulated as a MILP with a set of four constraints and a main binary sequencing variable.

These impose the following rules: only one model may be processed at any slot of the sequence, the demand for all model units shall be processed across all sequence slots, the processing of a model may only start if the processing of the previous one has already been completed, and the processing of a model is restricted to the time window of the station.

We expand this by adding a binary worker assignment variable and two constraints to regulate it. These impose that a worker can only be assigned to a workstation and that each workstation shall always have only one worker assigned to it. This worker assignment variable is added to the process time calculation in constraints 3 and 4 of the original problem.

To introduce competencies, we consider all operators to be able to do all tasks, but processing times vary depending on operator competence levels. Each task in a workstation has a competence requirement value. If both of these values are equal, then the task will be finished in its standard time. If an operator exceeds or fails to meet this requirement, it will work faster or slower than the task processing time, respectively. To guide the allocation process towards competence-efficient worker assignments, we include a competence gap variable, which is also featured in the objective function as a value to be minimized. Competence gaps are calculated in an additional constraint.

3 Preliminary results

We have conducted tests using a series of instances that feature seven workstations and workers and cycle times of around 30 minutes. Instances differ on the specific sequence, competency levels of the workforce, and cycle time.

As a result of the additional variables and constraints introduced, model complexity increases significantly. Runtimes markedly increase for all instances compared to MMSP-W. Nevertheless, the inclusion of worker assignment variables permits observing the influence of worker competence in the results. Compared to MMSP-W, instances with workforces with lower average competence levels show higher work overloads. The reverse is also true; workforces with higher competence levels incur in lower workloads.

This extension of the model can result in more accurate and realistic scheduling planning in manual MMALs. Despite this, we will explore the relationship between worker competence levels and processing times, as well as decompositions and heuristic methods in order to reduce the high solving times caused by the increase in model complexity.

References

[1] Boysen, N., Fliedner, M., & Scholl, A. (2009). Sequencing mixed-model assembly lines: Survey, classification and model critique. European Journal of Operational Research, 192(2), 349-373.

[2] De Bruecker, P., Van den Bergh, J., Beliën, J., & Demeulemeester, E. (2015). Workforce planning incorporating skills: State of the art. European Journal of Operational Research, 243(1), 1-16.

[3] Scholl, A., Klein, R., & Domschke, W. (1998). Pattern based vocabulary building for effectively sequencing mixed-model assembly lines. Journal of Heuristics, 4(4), 359-381.

A Comparative Study of Methods for Solving the Large Flexible Assembly Job Shop Scheduling Problems

L. Boveroux
University of Liège, Department of Electrical Engineering and Computer Science
laurie.boveroux@uliege.be

D. Ernst
University of Liège, Department of Electrical Engineering and Computer Science
dernst@uliege.be

Q. Louveaux
University of Liège, Department of Electrical Engineering and Computer Science
q.louveaux@uliege.be

Keywords: Assembly Flexible Job Shop

1 The Flexible Assembly Job Shop Scheduling Problem (FAJSP)

Scheduling is a decision-making process and plays an important role in diverse industries, including manufacturing, transportation and healthcare. It refers to the process of efficiently allocating a set of jobs to machines in order to minimise one or more objectives. In the context of Job Shop Scheduling Problems (JSPs), each job j consists of a set of operations $\{O_{j,1}, O_{j,2}, \ldots, O_{j,n}\}$ that must be performed in a given linear order on specific machines.

The Flexible JSP introduces flexibility by allowing operations to be processed on multiple machines. There exist alternative machines capable of performing the same operation. The assembly aspect introduces another dimension by defining the precedence between operations by a directed acyclic graph rather than a linear order. The goal of scheduling problems is to determine the optimal assignment of operations to machines and the sequencing of these operations while minimising an objective. The objective to be minimized in this paper is the sum of job completion times, reflecting the desire of industries to maximize billings. Working with instances of FAJSP has the goal of closely replicating real industrial scenarios.

2 Methods

This paper presents a comparative study that evaluates the effectiveness of seven methodologies in addressing the challenges posed by FAJSPs, particularly in the context of large-scale instances. The methods include:

1. Dispatching Rules (DPRs) [1],
2. Load Balancing Heuristics [1],
3. Genetic Algorithm (GA) [2],
4. Constraint Programming (CP),
5. Mixed Integer Linear Programming (MILP) [3],
6. Large Neighborhood Search (LNS),
7. Deep Reinforcement Learning (DRL) [4].

This comparative analysis aims to provide insights into the strengths and weaknesses of each method, specifically their performance on large-scale instances of the FAJSP created with the generator given in [3]. The results should facilitate decision-making by practitioners and researchers dealing with complex large-scale FAJSPs.

References

[1] Reijnen, R., van Straaten, K., Bukhsh, Z., & Zhang, Y. (2023). Job Shop Scheduling Benchmark: Environments and Instances for Learning and Non-learning Methods. arXiv preprint arXiv:2308.12794.

[2] Zhang, G., Gao, L., & Shi, Y. (2011). An effective genetic algorithm for the flexible job-shop scheduling problem. Expert Systems with Applications, 38(4), 3563-3573.

[3] Birgin, E. G., Feofiloff, P., Fernandes, C. G., De Melo, E. L., Oshiro, M. T., Ronconi, D. P. (2014). A MILP model for an extended version of the flexible job shop problem. Optimization Letters, 8, 1417-1431.

[4] Song, W., Chen, X., Li, Q., & Cao, Z. (2022). Flexible job-shop scheduling via graph neural network and deep reinforcement learning. IEEE Transactions on Industrial Informatics, 19(2), 1600-1610.

Strategic assembly line feeding decisions

H. Verplancke
Ghent University, Department of business informatics and operations research
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium
hannah.verplancke@ugent.be

V. Limère
Ghent University, Department of business informatics and operations research
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium
veronique.limere@ugent.be

E.-H. Aghezzaf
Ghent University, Department of Industrial Systems Engineering and Product Design
FlandersMake@UGent — core lab ISyE, Ghent, Belgium
elhoussaine.aghezzaf@ugent.be

E. Thanos
Ghent University, Department of business informatics and operations research
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium
emmanouil.thanos@ugent.be

Keywords: Assembly line feeding, Strategic process design

The ever-increasing needs of customers have pushed production companies to high mix low volume assembly line systems. Customized products have become the rule rather than an exception, which causes companies to produce large numbers of product variants on the same assembly line. This results in multiple challenges for production companies, such as the logistic difficulty of storing a growing number of parts and feeding them from the warehouse to the assembly line. While the number of parts is increasing, the available space at the border of line remains the same and is limited, not allowing to store full pallets of every individual stock keeping unit at its corresponding work place area. Therefore, different ways of stocking and feeding must be considered. The assembly line feeding problem studies ways to take up this intralogistic challenge in production environments.

Looking at the existing literature on the assembly line feeding problem, some relevant

solutions have come up in recent years. The introduction of supermarkets i.e., preparation areas where parts can be repacked, and five different part feeding policies helps companies to develop a tailored plan for every part. Diving deeper in the literature a firm focus on the tactical decision level can be noticed. The routing of both parts and vehicles and the assignment of feeding policies to parts have all been studied extensively. On the other hand, the strategic decision making on assembly line feeding, i.e., intralogistics design decisions, still remains largely unexplored.

These strategic capacity planning decisions mainly include facility lay-out and vehicle selection and sizing. Starting from an empty production site, companies have to decide on the number of supermarkets to construct, their locations and sizes. On top of that, the type and number of vehicles to be purchased must also be considered. In literature, those decisions have generally been studied separately, frequently resulting in restrictive assumptions.

The present study develops a mathematical model for the strategic assembly line feeding problem to fill this gap in the existing literature. The goal of the model is to minimize the total costs, which include the costs for transportation, replenishment, preparation, construction, extension and usage. To construct a solution applicable in real-life production environments, a number of constraints need to be incorporated in the model. Limited space at the border of line, non-overlapping restrictions for supermarkets and limited capacities for supermarkets, vehicles and kits are considered.

A Causal Machine Learning Approach for the Influence Maximization Problem

D. Caljon
KU Leuven, LIRIS
daan.caljon@kuleuven.be

J. Van Belle
KU Leuven, LIRIS
jente.vanbelle@kuleuven.be

W. Verbeke
KU Leuven, LIRIS
wouter.verbeke@kuleuven.be

Keywords: Spillover Effects, Causal Inference, Graph Neural Networks

Social networks play a vital role in word-of-mouth effects, product adoption, disease outbreak modeling, and more. A well-studied and important topic within the (social) network analysis literature is the problem of finding the set of entities to target for treatment to maximize the total effect of the treatment assignment, e.g., the set of customers to target in a marketing campaign to maximize the number of people who adopt a product. Kempe et al. (2003) formally introduced this problem as the Influence Maximization (IM) problem. The problem is defined as follows: given a graph G and a diffusion process, i.e., a model of how information or behavior propagates through the network, find the set of k seed nodes to maximize the total influence spread through the network.

There are some limitations to the current methods to model and solve the IM problem.

Node features. In the IM literature, node features are mostly ignored. Nevertheless, these features, such as age, might contain important information about how well information spreads through the network. For example, young people might share more about what they (want to) buy with their friends than older people. Incorporating these features in the optimization procedure should improve the treatment assignment.

Individualized degree of susceptibility. To the best of our knowledge, all methods assume that each node that is selected for treatment, will also be activated, e.g., buy the product. However, in reality, some people might not be susceptible to a certain campaign. These customers are called lost causes and sure things in the uplift modeling literature. Lost causes will never buy the product, whereas sure things will always buy the product, regardless of the treatment. This highlights the existence of an individual treatment effect (ITE): individuals react differently to a certain treatment depending on their characteristics. One might only want to target those customers that react positively to the treatment while trying to avoid customers that do not react, or react negatively, e.g., cus-

tomers targeted by a retention campaign might get triggered to churn (Devriendt et al., 2018).

Diffusion processes are model-based. The most widely used diffusion processes to solve the IM problem often rely on assumptions rather than on data. However, solving the IM problem with a misspecified diffusion model might give widely different solutions than the optimal seed set. Consequently, if the diffusion process is not empirically valid, the found solution using one of the "classic" model-based diffusion processes might lead to undesired outcomes (Aral and Dhillon, 2018).

Our solution. To address these limitations, we propose a causal machine learning solution for data-driven estimation of a node's susceptibility and spillover effect using node features. The spillover effect refers to the effect of giving treatment to one node, i.e., including it in the seed set, on other nodes' outcomes (Ogburn and VanderWeele, 2014). For this purpose, we use a Graph Neural Network (GNN) based causal estimator inspired by NetEst (Jiang and Sun, 2022). In this work, we show that using a two-step approach, i.e., training a causal estimator and using its predictions to optimize the treatment allocation, improves the expected influence spread compared to traditional IM and uplift modeling approaches. This improvement is observed in semi-synthetic experiments using both simulated and real networks. To optimize the treatment allocation, we employ modified versions of classic IM algorithms such as the greedy algorithm proposed by Kempe et al. (2003) and the genetic algorithm proposed by Bucur and Iacca (2016).

References

Aral, S. and Dhillon, P. S. (2018). Social influence maximization under empirical influence models. *Nature human behaviour*, 2(6):375–382.

Bucur, D. and Iacca, G. (2016). Influence maximization in social networks with genetic algorithms. In *Applications of Evolutionary Computation: 19th European Conference, EvoApplications 2016, Porto, Portugal, March 30–April 1, 2016, Proceedings, Part I 19*, pages 379–392. Springer.

Devriendt, F., Moldovan, D., and Verbeke, W. (2018). A literature survey and experimental evaluation of the state-of-the-art in uplift modeling: A stepping stone toward the development of prescriptive analytics. *Big data*, 6(1):13–41.

Jiang, S. and Sun, Y. (2022). Estimating causal effects on networked observational data via representation learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 852–861.

Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146.

Ogburn, E. L. and VanderWeele, T. J. (2014). Causal diagrams for interference. *Statistical Science*, 29(4):559–578.

Realistic Model Benchmarking using Synthetic Data for Prescriptive Process Monitoring

J. De Moor
KU Leuven, LIRIS
jakob.demoor@kuleuven.be

H. Weytjens
KU Leuven, LIRIS
Technical University of Munich, CIT
hans.weytjens@{kuleuven.be, tum.de}

J. De Smedt
KU Leuven, LIRIS
johannes.desmedt@kuleuven.be

J. De Weerdt
KU Leuven, LIRIS
jochen.deweerdt@kuleuven.be

Keywords: Synthetic data, Causal Inference, Reinforcement Learning

Prescriptive Process Monitoring. The emerging field of Prescriptive Process Monitoring (PresPM) focuses on devising methods to suggest interventions in business processes with the goal of optimizing them and attaining desired outcomes [1, 2].

Interventions in PresPM can be categorized across three dimensions, illustrated here through a bank's loan application process. First of all, the dimension data impact distinguishes interventions into those that influence the sequence of executed process activities (e.g., a cancellation of the application) – known as the control flow – and those that influence only a data attribute (e.g., changing the interest rate). The dimension action width pertains to the range of possible actions available during an intervention [3]. An intervention can be categorized as binary (e.g., cancel the application or not), multi-class (e.g., choosing to call, email or visit the applicant) or continuous (e.g., when setting the interest rate, options form a continuous range). Action depth, the third dimension, characterizes the extent to which intervening actions can be executed consecutively, along with the timing of these actions [3]. In this dimension, interventions can be classified into two perspectives: those with fixed timing in the process (e.g., choosing between a fast and slow procedure once the application is submitted) and those without (e.g. doing a background check), as well as those involving a chain of intervening actions and those that do not.

Ideally, a PresPM model should be capable to generate pertinent intervention suggestions for each individual process case. However, given the causal relationship between an intervention and a case outcome, evaluating models using historical data is challenging because of the absence of ground truths [3]. While historical data is necessary for a (partial) indication of a model's real-world performance, assessing a model's policy on an individual case basis is infeasible because the intervention is either present or absent in

a given case. The concept of a counterfactual for a case involves considering what would have occurred if an intervention had been applied when it was originally absent in the given case, or vice versa. While having access to these counterfactuals would address the evaluation problem, such accessibility is not viable when working with historical data. Consequently, synthetic data can be effective for model evaluation, facilitating the generation of process cases and their counterfactuals. Another advantage of synthetic data is the flexibility it offers in manipulating the data generation. The generative functions are known and can be easily adjusted to setup different experiments focusing on a variety of problems (e.g., intervention types and combinations of them).

Research gaps. A first research gap concerns the lack of comprehensive synthetic data. Despite their significant potential for PresPM, synthetic datasets often suffer from issues of being unrealistic, oversimplified, and predominantly tailored to singular types of interventions. Secondly, studies typically involve a limited array of techniques and interventions considered concurrently. As a result, both research gaps present opportunities for improvement to enable comprehensive comparisons of machine learning techniques within PresPM.

Proposed approach. We propose an extensive comparison of machine learning techniques, employing a synthetic dataset with the aim of covering as many intervention dimensions as possible. This dataset is designed to emulate a loan application process within a local bank unit, striking a balance between real-life complexity and simplicity essential for drawing meaningful conclusions in method comparison. Notably, the dataset incorporates a diverse array of interventions across each dimension, facilitating method selection for various intervention types. We structure the data generation process to allow for straightforward modification of the data generative functions.

Our study employs diverse machine learning approaches, assuming an unknown initial process model. We compare variants of Causal Inference and (offline/online) Reinforcement Learning, using the synthetic data first to easily obtain clear insights. Whenever feasible, we validate the reliability of these insights in practical settings by employing historical datasets, leveraging the aforementioned advantage of historical data to (partially) indicate real-world performance.

In summary, our contribution involves conducting comprehensive and realistic model benchmarking. We achieve this by utilizing a flexible data generator, considering various intervention types, and employing multiple machine learning techniques.

References

[1] Kateryna Kubrak, Fredrik Milani, Alexander Nolte, and Marlon Dumas. Prescriptive process monitoring: Quo vadis? *CoRR*, abs/2112.01769, 2021.

[2] Mahmoud Shoush and Marlon Dumas. When to intervene? prescriptive process monitoring under uncertainty and resource constraints. In Claudio Di Ciccio, Remco Dijkman, Adela del Río Ortega, and Stefanie Rinderle-Ma, editors, *Business Process Management Forum*, pages 207–223, Cham, 2022. Springer International Publishing.

[3] Hans Weytjens, Wouter Verbeke, and Jochen De Weerdt. Timing process interventions with causal inference and reinforcement learning. *arXiv preprint arXiv:2306.04299*, 2023.

CurveMMoE: A Deep Multi-Task Learning with Multi-Gate Mixture of Experts Approach for Trading Commodity Futures Curves

J. Ong
Singapore University of Technology and Design, Information Systems Technology and Design
joel_ong@mymail.sutd.edu.sg

D. Herremans
Singapore University of Technology and Design, Information Systems Technology and Design
dorien_herremans@sutd.edu.sg

Keywords: Deep Multi-Task Learning (MTL), Multi-Gate Mixture of Experts (MMoE), Commodity Futures Curve Trading

1 Introduction

Commodity futures curve trading is a sophisticated financial strategy that involves analyzing and speculating the price movement of commodity futures curve over timer [1]. This type of trading is grounded in understanding the futures curve, which represents the prices of futures contracts for commodities at different expiration dates. At the core of this strategy is the concept of the futures curve. There are two primary shapes that a futures curve can take: contango and backwardation. Contango occurs when futures prices are higher in the forthcoming months than in the immediate months. Conversely, backwardation is when future prices are lower than the current prices.

By understanding the relationships between different contract expiration dates, a Trader can leverage these relationships to generate profits. Traders may employ various tools, such as technical analysis and fundamental analysis, to predict the curve's movement more accurately. Increasingly, Traders are also turning to advanced machine learning and deep learning models. These models provide an opportunity to understand better and predict the intricacies involved in the dynamics of futures curve movements.

2 Methodology

This work introduces CurveMMoE, a novel deep-learning framework for strategic commodity futures curve trading. Utilizing a Deep Multi-Task Learning (MTL) approach,

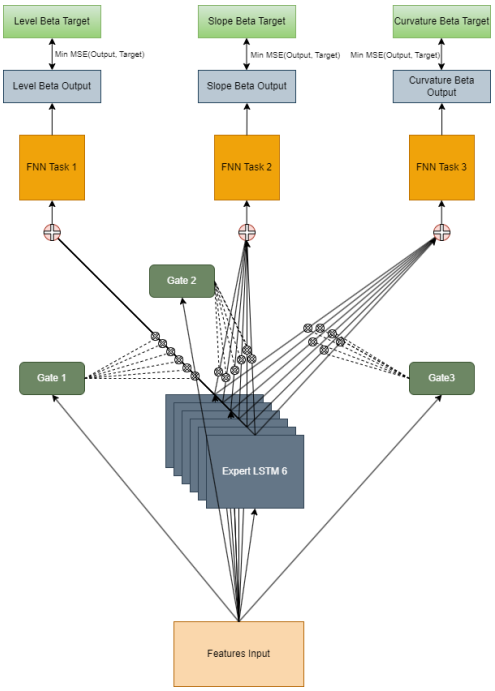

Figure 1: The figure depicts our proposed Multi-task learning architecture with Multi-Gate Mixture of Experts. The architecture consists of 6 Long-Short Term Memory (LSTM) modules, each called an expert. Each task-specific layer is a feedforward network (FNN) layer, whereby each FNN learns a set of parameters optimized to their respective task. In addition, Gating networks are introduced for each task; gating networks assemble the experts with different weights, allowing different tasks to utilize the experts differently.

CurveMMoE incorporates a Multi-Gate Mixture of Experts (MMoE) model [2], allowing for the simultaneous execution of various predictive tasks in the domain of futures curve trading. The essence of CurveMMoE lies in its MMoE model, which coordinates a range of specialized sub-models, each responsible for different aspects of market data analysis, thus enhancing the system's overall predictive capability.

3 Conclusion

In our comprehensive evaluation using diverse commodity futures datasets, CurveMMoE exhibits significant improvements in predictive accuracy and trading efficacy compared to benchmark models. CurveMMoE establishes itself as the state-of-the-art approach for trading commodity futures curves. It achieves this through a sophisticated blend of Deep Multi-Task Learning and integrating a diverse array of expert models. The insights gained from this research have broader implications, shedding light on the potential of artificial intelligence to revolutionize financial market analysis.

References

[1] R. J. Bianchi et al. "Exploiting the dynamics of Commodity Futures Curves". In: *Journal of Banking & Finance* 154 (2023), p. 106965. DOI: 10.1016/j.jbankfin.2023.106965. URL: https://doi.org/10.1016/j.jbankfin.2023.106965.

[2] Jiaqi Ma et al. "Modeling Task Relationships in Multi-Task Learning with Multi-Gate Mixture-of-Experts". In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '18. London, United Kingdom: Association for Computing Machinery, 2018, pp. 1930–1939. ISBN: 9781450355520. DOI: 10.1145/3219819.3220007. URL: https://doi.org/10.1145/3219819.3220007.

Counterfactual Explanations for Predictive Process Analytics

A. Stevens
KU Leuven, FEB
Leuven, Belgium
alexander.stevens@kuleuven.be

C. Ouyang
QUT, School of Information Systems
Brisbane, Australia

J. De Smedt
KU Leuven, FEB
Leuven, Belgium

C. Moreira
UTS, Human Technology Institute
Sydney, Australia

Keywords: Counterfactuals, Manifold Learning, Declare Templates

Process mining, situated at the intersection of data science and process management, extracts insights from executed cases of processes within information systems. The focus of Predictive Process Analytics (PPA), a subfield of process mining, is on exploring process-related trends such as identifying bottlenecks (e.g., how long will it take to process my healthcare procedure?), future activities (e.g., will a specific diagnostic test be necessary for this case?) or predicting outcomes (e.g., will my patient be admitted to the Intensive Care Unit (ICU) or not?). In the latter case, the challenge lies in the ability to predict the future state from observed data - an *event* we aim to predict but has not occurred yet. These insights help to enhance operational efficiency, reduce costs, increase productivity, and enhance customer satisfaction across diverse industries such as healthcare, finance, logistics, and manufacturing [2].

In recent years, the advent of sophisticated deep learning architectures has shown the potential of *black-box* models in predicting the outcome of uncompleted cases. However, their extensive parametric space with hundreds of layers and parameters results in a lack of transparency. This opacity impedes users from comprehending the rationale behind a particular prediction, prompting the need for Explainable Artificial Intelligence (XAI). Given the AI transparency regulations outlined in the General Data Protection Regulation (GPDR) and the proposed European Union regulation for Artificial Intelligence (AI), known as the AI Act, counterfactual explanations emerge as a highly promising approach for providing human-understandable explanations [3]. While factual explanations offer a more distinct understanding of the relative importance of each feature in influencing the model output, a ”what if” analysis offers users a set of changes that an individual should undertake to avoid certain undesired outcomes.

Counterfactual explanations are highly relevant in the field of PPA to help reasoning about (and preventing) negative outcomes such as patient ICU admissions, equipment failures, process deviations, and quality issues. For example, in a scenario where a black box model predicts whether a patient will be admitted to the ICU (or not), a counterfactual explanation would reveal the changes in the input necessary to prevent the predicted admission to the ICU. However, not all counterfactuals may propose changes or interventions that are feasible or realistic in the given context [3]. Therefore, we adapt the most important properties that constitute a valid counterfactual [1] such as feasibility, plausibility, proximity, diversity and sparsity, to make them suitable for a process-based analysis.

In this paper, we introduce a novel approach for obtaining *feasible* and *plausible* counterfactual explanations for predictive process analytics. First, we use Declare language templates to learn sequential patterns in the cases. Patterns that are present in all cases, i.e. *trace* Declare patterns, define the minimal requirement for a case, while *label-specific* Declare patterns provide insight into the behavioural characteristics shared among cases with the same outcome. To enforce *plausibility*, we enforce that our tailored approach does not violate any label-specific Declare patterns when generating counterfactuals. Ensuring the feasibility of our approach involves constraining the algorithm to generate scenarios within high-density regions of the process data, i.e. aligned with the observed data distribution. We additionally implement that our algorithm does not violate any trace Declare patterns to ensure alignment with the historical cases.

Finally, We present a case study in which we assess the applicability and usefulness of our counterfactual generation algorithm, and we introduce a post-hoc evaluation framework to assess other important properties such as *proximity*, *sparsity* and *diversity*.

References

[1] Yu-Liang Chou, Catarina Moreira, Peter Bruza, Chun Ouyang, and Joaquim Jorge. Counterfactuals and causability in explainable artificial intelligence: Theory, algorithms, and applications. *Information Fusion*, 81:59–83, 2022.

[2] Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A Reijers. Introduction to business process management. *Fundamentals of business process management*, pages 1–33, 2018.

[3] Catarina Moreira, Yu-Liang Chou, Chihcheng Hsieh, Chun Ouyang, Joaquim Jorge, and João Madeiras Pereira. Benchmarking counterfactual algorithms for xai: From white box to black box. *arXiv preprint arXiv:2203.02399*, 2022.

An Efficient Decomposition Matheuristic for the Unit Commitment Problem with Power Stability Constraints

E. M. Er Raqabi[1,2], A. Bani[1,2]
[1] Polytechnique Montreal & GERAD, Quebec, Canada,
Department of Mathematics and Industrial Engineering

M. Morabit[2], A. Blondin[2],
A. Besner[2], I. El Hallaoui[1]
[2] Hydro Quebec Research Institute (IREQ), Quebec, Canada,
Data Science and High Performance Computing

Keywords: MILP, Optimization, Matheuristic, Unit Commitment.

1 Trigger

Hydro-Quebec (HQ), a vertically integrated utility, produces, transmits, and distributes the majority of the power in the province of Quebec. The vast 735kV transmission line and significant hydroelectric dams in the far north give the power grid a unique architecture, enabling the generation to satisfy the total demand, located thousands of kilometers in southern Quebec. Compelled by the grid's uniqueness, HQ created significant procedures for tracking the system's stability. Several stability phenomena (e.g., voltage, frequency, and angular stabilities) are the sources of those stability limits. These limits are non-linear. Still, HQ must incorporate them in the design and generation operation.

2 Contributions

We tackle HQ's unit commitment (UC) problem with power stability constraints (Valette et al., 2009). After linearizing these constraints, we present the mixed-integer linear programming (MILP) model, which incorporates power stability restrictions into HQ's UC and reserve monitoring tools. Then, we conduct an investigation strategy to identify complexity sources. After that, we develop an efficient matheuristic to solve the MILP. We highlight computational results on real-life instances.

The contribution is fourfold:

1. An exploratory analysis to identify the complexity sources of the UC problem with power stability constraints.
2. An efficient decomposition matheuristic with various variants to solve the MILP model.
3. A computational study on real-world instances to show the performance of the suggested matheuristic.
4. A post-computational analysis to highlight several managerial insights.

3 References

[1] Valette, A., Huang, J. A., Guillon, S., Loud, L., Vanier, G., Levesque, F., Riverin, L., Rizzi, J.C. & Guillemette, F. (2009). An integrated approach for optimizing dynamic transfer limits at Hydro-Quebec. IEEE Transactions on Power Systems, 24(3), 1310-1317.

A Matheuristic for the Optimal Placement of EV Charging Stations in an Urban Environment

B. Coulier
KU Leuven, VITO, EnergyVille, Department of Computer Science
bryan.coulier@kuleuven.be

H. Çalık
KU Leuven, EnergyVille, ELECTA

T. Becker
VITO, EnergyVille, AMO

G. Vanden Berghe
KU Leuven, Department of Computer Science

Keywords: EV Location Optimization, Power Grid Optimization, Grid Decomposition

In response to global environmental concerns and policy initiatives, the automotive industry is currently undergoing a significant transformation towards sustainable transportation. The surge in Electric Vehicle (EV) adoption is one of the most identifiable and fast-moving components of this transition. Moreover, this particular shift is set only to intensify further in the coming years with governments worldwide setting ambitious targets to reduce CO_2 emissions and phase out Internal Combustion Vehicle (ICV) sales [2].

As EVs become a major component of the transition to a greener, more electrified future, governments are investing heavily in charging infrastructure. The rapid rise of EV adoption poses significant challenges to conventional power grids, particularly transformer overloading and voltage fluctuations arising from the proliferation of charging stations [1]. Intelligent solutions are necessary in order to ensure reliable and efficient charging infrastructure.

Current studies concerning EV charging infrastructure often overlook the underlying power distribution network and the accessibility of charging stations for EV users. The majority of these studies are small in scale, employing metaheuristics without validating their results using exact solutions. Moreover, researchers have predominantly focused on fast chargers, while slow chargers play an equally crucial role in urban settings. Slow charging stations are cost-effective, require minimal grid investments, and are conducive to sustainable day-to-day charging.

In an effort to address these significant gaps, our research focuses on the placement of slow charging stations in urban areas while taking into account (i) the distribution net-

work, (ii) customer demand, and (iii) accessibility. We propose a novel Mixed-Integer Linear Programming (MILP) model to determine the optimal placement, quantity, and type of charging stations, considering both charging station installation costs and power generation costs. To enhance scalability and computational speed, we also introduce a multi-stage matheuristic for solving large-scale instances by decomposing the grid network.

The proposed matheuristic offers a robust solution method for optimizing the placement of slow charging stations in urban areas, ensuring reliable EV charging while minimizing grid impact.

References

[1] F. Ahmad, A. Iqbal, I. Ashraf, M. Marzband, et al. Optimal location of electric vehicle charging station and its impact on distribution network: A review. *Energy Reports*, 8:2314–2333, 2022.

[2] E. Commission. European green deal. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_nl, 2023. Accessed: 01-08-2023.

A mixed-integer-programming formulation for the traveling student problem

Q. Louveaux
University of Liège, Computer Science and Electrical Engineering Department
q.louveaux@uliege.be

Keywords: Room allocation, mixed-integer programming

1 Abstract

We consider the problem of allocating rooms to classes assuming that the timetable of courses of a school or university is fixed. In this context, the first set of hard constraints is to assign a sufficiently large room in order to accomodate all students of a given class. We define the traveling student problem as the problem of allocating the rooms in such a way that the students have the minimum distance to travel during a course day. In general, this problem can be formulated as a quadratic assignment problem which is too expensive to solve for problems with a size typical of those faced by universities. In this talk, we propose to relax the problem in order to take into account the large distances that a student may face and to minimize specifically this type of events. We then propose a mixed-integer programming formulation for this relaxation. Computational experiments show that this approach is successful in order to tackle instances of real-life size. In particular, our data set comes from allocating rooms for the courses at the Liège campus of the University of Liège.

Quantum Computing in Logistics and Supply Chain Management - an Overview

F. Phillipson
Maastricht University, School of Business and Economics
TNO, Applied Cryptography and Quantum Algorithms
The Netherlands - f.phillipson@maastrichtuniversity.nl

Keywords: Quantum Computing, Optimisation, Logistics and Supply Chains

1 Introduction

The work explores the integration of quantum computing into logistics and supply chain management, highlighting its potential for solving complex optimisation problems. The challenges in this field involve route planning, inventory management, production scheduling, and more. Addressing conflicting objectives and dynamic system behaviour adds complexity to decision-making processes. Advanced techniques, such as mathematical modelling and metaheuristics, are commonly employed from the field of Operations Research to solve these problems.

Quantum computing, a transformative paradigm, leverages principles from quantum mechanics to surpass classical computers' capabilities. Advancements include improved database searching, integer factorisation and, expectedly, combinatorial optimisation. The latter can be found in logistics and supply chain management, including vehicle routing, facility location, and scheduling problems.

2 Current State and Challenges

While quantum computing offers great promise, practical, large-scale implementations are in the early stages. However, quantum hardware is evolving rapidly, and researchers are developing software stacks and algorithms for this hardware. The methodology used involved an extensive literature search using Google Scholar and Scopus, focusing on quantum algorithms and logistic problem categories. Our work categorises logistics and supply chain problems into six pivotal categories: routing, logistic network design, fleet maintenance, cargo loading, prediction, and scheduling. It provides an overview of quantum approaches used to provide approximate solutions to these problems. The overview of around 70 published papers in the field of logistics and supply chain optimisation highlights ongoing research worldwide.

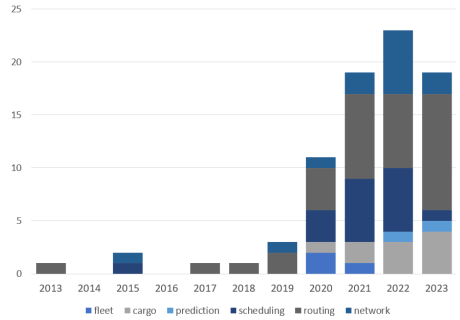

Figure 1: Distribution of the papers on quantum approaches for logistics and supply chain management over the years, per category: **fleet** optimisation, **cargo** loading, **prediction**, **scheduling routing** and **network** design

3 Conclusions and Recommendations

The literature review underscores a prevalence of hybrid quantum solutions, aligning with two key considerations. Firstly, the limited size of current hardware. Secondly, for many optimisation problems, a hybrid approach is anticipated to persist as quantum computers are not expected to solve NP-hard problems [1].

Another notable observation in the overview is the widespread use of Quantum Annealing as the primary quantum paradigm, orginating from the technology's higher qubit count. Gate-based devices, including Quantum Approximate Optimisation Algorithm (QAOA) and other Variational Quantum Eigensolver (VQE) approaches, rank as the second most adopted approach. The landscape still offers significant room for experimentation, particularly within more problem-specific and machine learning-oriented approaches. An examination of covered topics reveals a predominant focus on classical areas such as routing and scheduling. Notably, the prediction domain emerges as an area with untapped potential for further research, particularly in the realm of Quantum Machine Learning (QML) approaches.

In general, our findings align with the conclusions drawn in [2]. Extending beyond the scope of Vehicle Routing Problem (VRP) and Traveling Salesman Problem (TSP) studies, our overview concurs that existing hardware is in its infancy, marked by a small number of qubits and inherent limitations. Researchers encounter challenges tailoring problems to hardware capacities, limiting their ability to address complex formulations effectively. While a minority of papers suggests the superiority of hybrid approaches over current classical methods, the majority underscore the ***potential*** advantages and benefits inherent in their proposed quantum solutions.

4 References

[1] Aaronson, S. (2008). The limits of quantum. Scientific American 298(3), 62–69
Osaba, E. et al. (2022): A systematic literature review of quantum computing for routing problems. IEEE Access 10, 55805–55817

A backtracking algorithm for finding bi-regular cage graphs

T. Van den Eede
KU Leuven, Department of Computer Science
tibo.vandeneede@student.kuleuven.be

J. Jooken
KU Leuven KULAK, Department of Computer Science
jorik.jooken@kuleuven.be

J. Goedgebeur
KU Leuven KULAK, Department of Computer Science
jan.goedgebeur@kuleuven.be

Keywords: Graph algorithms, Cage Problem, Extremal problems

The vertex degrees and cycles are one of the most fundamental notions of a graph. In this work, we focus on (r, g)-*graphs*, which are defined as graphs where all vertices have degree r and the shortest cycle is g edges long. It is known that (r, g)-graphs exist for all integers $r \geq 2$ and $g \geq 3$ [1], but much less is known about the smallest such graphs. These smallest graphs are called (r, g)-*cages*, which are (r, g)-graphs of minimum order. The order of an (r, g)-cage is denoted by $n(r, g)$. Figure 1 illustrates the only $(3, 6)$-cage, also known as the Heawood graph. Since this cage has 14 vertices, $n(3, 6)$ is equal to 14.

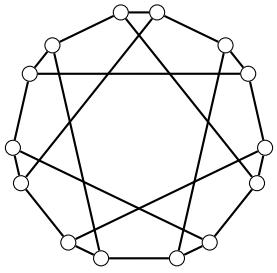

Figure 1: The Heawood graph.

The problem of determining $n(r, g)$, also known as the *Cage Problem*, is a notoriously difficult problem. A recent trend is to look at variants of (r, g)-cages in hopes of shedding more light on the Cage Problem. One such variant is the $(\{r, m\}; g)$-cage, also denoted

as the bi-regular cage. An $(\{r, m\}; g)$-*cage* is a graph of minimum order with girth g and vertices of degree r and m, with $r < m$.

In the history of the field, backtracking algorithms have been able to find unknown values for $n(r, g)$ [2, 3]. These backtracking algorithms exhaustively generate all possible (r, g)-graphs of a given order based on the development of interesting pruning rules and appropriate modellings of the search space. In analogy to these algorithms, we will present a foundation for a backtracking algorithm to exhaustively generate $(\{r, m\}; g)$-graphs.

References

[1] H. Sachs (1963). Regular graphs with given girth and restricted circuits *J. London Math. Soc.* 38, 423 – 429.

[2] B. D. McKay, W. Myrvold, J. Nadon (1998). Fast Backtracking Principles Applied to Find New Cages. *9th Annual ACM-SIAM Symposium on Discrete Algorithms* 188 – 191.

[3] G. Exoo, B. D. McKay, W. Myrvold, J. Nadon (2011). Computational determination of (3,11) and (4,7) cages. *Journal of Discrete Algorithms*, Volume 9, Issue 2, 166 – 169.

On the injective chromatic number of a graph

M. Daneels
KU Leuven Kulak, Department of Computer Science
matias.daneels@student.kuleuven.be

J. Goedgebeur
KU Leuven Kulak, Departement of Computer Science
jan.goedgebeur@kuleuven.be

J. Renders
KU Leuven Kulak, Departement of Computer Science
jarne.renders@kuleuven.be

Keywords: Injective Chromatic Number, Algorithm, Backtracking

A k-*coloring* of a graph $G = (V, E)$ is a mapping $c : V \to 1, 2, ..., k$. The coloring c is *injective* if, for every vertex $v \in V$, all neighbours of v are assigned with distinct colors. The *injective chromatic number* $\chi_i(G)$ of G is the smallest k such that G has an injective k-coloring. Graph coloring has many applications in scheduling problems. In 2012, Chen et al. conjectured an upper bound for the injective chromatic number of planar graphs [1], stating that for each planar graph G, $\chi_i(G) \leq \lceil \frac{3}{2}\Delta \rceil$. This conjecture was disproved by Luzăr and Škrekovski who provided some counterexamples [3]. Luzăr and Škrekovski proposed the following conjecture for the injective chromatic number of planar graphs: Let G be a planar graph with maximum degree Δ. Then

(a) $\chi_i(G) \leq 5$, if $\Delta = 3$;

(b) $\chi_i(G) \leq \Delta + 5$, if $4 \leq \Delta \leq 7$;

(c) $\chi_i(G) \leq \lfloor \frac{3}{2}\Delta \rfloor + 1$, if $\Delta \geq 8$.

This conjecture, if true, has been proven to be sharp. The conjecture is closely related to the conjecture of Wegner on *square coloring* [4]. A coloring c is a square coloring if, for every vertex $v \in V$, all neighbours of v and v itself are assigned with distinct colors. If the conjecture of Wegner is true, the conjecture of Luzăr and Škrekovski is also true.

We study this topic with the use of an exact backtracking algorithm to calculate the injective chromatic number of graphs. For $\Delta = 3$, the conjecture of Luzăr and Škrekovski has been proven correct for *fullerene graphs* [2], which are planar cubic graphs that only contain pentagons and hexagons. They proved the conjecture in de case of *exact square coloring*. A coloring c is an exact square coloring if, for every vertex $v \in V$, every pair

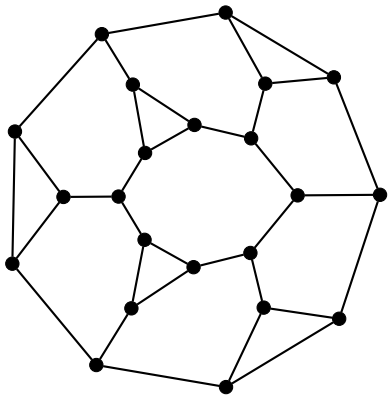

Figure 1: The second generation of an infinite family with $\Delta = 3$ and $\chi_i(G) = 5$

of neighbours of v need distinct colors if they are not neighbours of each other. If a graph does not contain any triangles, the exact square chromatic number is equal to the injective chromatic number. We found an infinite family of graphs with $\Delta = 3$ and $\chi_i(G) = 5$ that contains triangles (see Figure 1). We also found a way to, in some cases, be able to extend a graph where $\chi_i(G) \geq 9$ without changing $\chi_i(G)$. This provided us a way to make multiple infinite families with $\chi_i(G) = 14$ where $\Delta = 9$.

References

[1] Min Chen, Geňa Hahn, André Raspaud, and Weifan Wang. Some results on the injective chromatic number of graphs. *Journal of Combinatorial Optimization*, 24(3):299–318, 2012.

[2] Florent Foucaud, Hervé Hocquard, Suchismita Mishra, Narayanan Narayanan, Reza Naserasr, Eric Sopena, and Petru Valicov. Exact square coloring of subcubic planar graphs. *Discrete Applied Mathematics*, 293:74–89, 2021.

[3] Borut Lužar and Riste Škrekovski. Counterexamples to a conjecture on injective colorings. *Ars Mathematica Contemporanea*, 8(2), 2015.

[4] Gerd Wegner. Graphs with given diameter and a coloring problem. Technical Report, University of Dortmund, 1977.

Tuning house network graphs

M. Geerts
KU Leuven, LIRIS
margot.geerts@kuleuven.be

S. vanden Broucke
UGent, Business Informatics

J. De Weerdt
KU Leuven, LIRIS

Keywords: house prices, graph neural networks, graph construction

Recent advances in Deep Learning continue to prove the impressive predictive power of Graph Neural Networks (GNNs) to leverage graph structured data. A common use case is a citation network where nodes represent documents that are connected if one cites the other. While house price data sets are tabular data, the geographic location of houses provides the intuition to model housing data as a graph. That is, a graph structure allows to spatially position houses relative to each other. Few attempts have been made to predict house prices with graph neural networks [1, 3, 4, 5]. However, one of the main points of discrepancy within the literature is the construction of the housing network. Houses are connected based on geographic distance [1, 5], feature similarity [3], or through other nodes that represent their characteristics [4]. Distance and similarity metrics for edge construction require an additional choice, that is, a threshold to determine which houses to connect. Nevertheless, methodologies for setting this threshold and alternative network construction methodologies remain unreported.

We systematically investigate the construction of the housing network with respect to downstream price prediction performance in GraphSAGE models [2]. These models fit the house price prediction case well as they are inductive and generally perform well on large graphs. First, we provide a theoretical approach to optimal network architecture for GraphSAGE-based price prediction. Starting from a common distance-based housing network, we investigate edge removal and edge addition for optimizing this network. A data set of real estate transactions in Belgium is used to present the theoretical bound for GraphSAGE-based predictive performance. We show that an optimized housing network informed by the target can result in a considerable boost in performance of GraphSAGE models, leading to outperform traditional baselines such as Gradient Boosting. Second, we extend this theoretical approach towards inductive strategies, which replicate this theoretical approach performance-wise in a real-world practical setting.

References

[1] Sarkar Snigdha Sarathi Das, Mohammed Eunus Ali, Yuan-Fang Li, Yong-Bin Kang, and Timos Sellis. Boosting house price predictions using geo-spatial network embedding. *Data Mining and Knowledge Discovery*, 9 2021.

[2] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

[3] Konstantin Klemmer, Nathan Safir, and Daniel B. Neill. Positional encoder graph neural networks for geographic data. In *International Conference on Artificial Intelligence and Statistics*, pages 1379–1389. PMLR, 2023.

[4] Hao Peng, Jianxin Li, Zheng Wang, Renyu Yang, Mingsheng Liu, Mingming Zhang, Philip Yu, and Lifang He. Lifelong property price prediction: A case study for the toronto real estate market. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1, 2021.

[5] Di Zhu, Yu Liu, Xin Yao, and Manfred M Fischer. Spatial regression graph convolutional neural networks: A deep learning paradigm for spatial multivariate distributions. *GeoInformatica*, pages 1–32, 2021.

Graphical Processing for Fast Survivability Assessments

T. Vancaeyzeele
Royal Military Academy, Department of Weapon Systems and Ballistics
tom.vancaeyzeele@dymasec.be

J. Gallant
Royal Military Academy, Department of Weapon Systems and Ballistics

B. Lauwens
Royal Military Academy, Department of Mathematics

Keywords: Filters, Survivability, Graphical Processing, Cartographic Projection, Intercept Envelope

1 Abstract

One of the current methods to estimate or predict the probability of survival for a new type of threat and to evaluate the effectiveness of a new effector involves ray-tracing and Monte Carlo Simulation based on the predicted fragmentation program. During the design phase of a fragmentation warhead, the predicted average fragment distribution pattern including the ejection angle and ejection velocity can be estimated using a dedicated software such as Split-X. Once the fragment information is available, the effects are evaluated by tracing the individual fragment trajectories. Here, the trajectories are considered to be linear, ignoring the different forces acting upon the fragment during its flight. This approximation can only be used when considering short flight times.

With a 3D model of the threat available, the effects of the warhead can be simulated. Here the individual fragment trajectories are calculated as rays and their interaction with the 3D model is registered. An interaction occurs when the ray passes through the structure, in general modeled as a mesh structure. Through this calculation, the scenario-specific probability of hitting the threat can be determined. To obtain the average probability, the simulation must be run a large number of times varying the initial roll angle, until a convergence has been attained. By randomly changing the roll angle, the uncertainty of the final orientation can be simulated.

The projectiles are in general loaded with a random initial roll angle for the warhead. The spin damping and acceleration which is partly determined by the muzzle velocity, barrel twist, axial moment of inertia and aerodynamic parameters is stochastic in nature. It is reasonable to expect the final roll angle to follow a uniform probability distribution. For the ejection of the fragments, using the Taylor angles and explosives behaviour equations will lead to a velocity vector for each fragment. Without further indications, this vector can be assumed as the average fragment behaviour, which in turn can be defined using a normal distribution.

At close ranges, for slow threats, the speed of the fragments could be replaced with the relative speed (with respect to the threat) and the travel time for the fragments can be ignored. Considering the warhead as the centre of the reference system, the threat can be projected on a virtual cylinder through a Mercator or cylindrical projection. When using a black and white representation of the projected threat, the positional information of the threat is translated to a 2D projection. For the fragments, this would allow reducing the problem further by converting the fragment distribution pattern to a one dimensional filter aligned with the projectile axis, on the projection surface. Using convolution, the predicted probability of hitting the threat can be determined by analysing the projection image after convolution. Both probability of hit information and the average number of predicted hits are available in the resulting image.

This approach speeds up probability of hit estimations significantly, which allows faster estimations of the effectiveness of a certain effector against a ”new” type of threat, allowing more time for the Fire Control Unit to determine the optimal firing solution.

This method also greatly increases the calculation speed when considering multiple threats in a fleet or swarm, in which case determining a good defence strategy is more complex. This can be accomplished, for example, by using different colours or masks while generating the threat projections. Through the use of the ”visualisation buffer”, targets or areas of targets which are ”obstructed” by obstacles or other threats are easily removed from the final probability of hit analysis by using masks. This way, fragments hitting a threat are not necessarily also considered hitting other ”aligned” threats further along the undisturbed, ideal fragment trajectory.

This method proves it can provide good probability of hit approximations a lot faster, using less computing time and capacity, leaving more time for the Fire Control Computer to search for better firing solutions under uncertain conditions.

A research framework for algorithm engineering

B. Depaire
Hasselt University, Research Group Business Informatics
benoit.depaire@uhasselt.be

J. Mendling
Humboldt-Universität zu Berlin, Wirtschaftsuniversität Wien, Weizenbaum-Institut
jan.mendling@hu-berlin.de

H. Leopold
Kühne Logistics University, Hasso Plattner Institute, University of Potsdam
Henrik.Leopold@klu.org

H. Meyerhenke
Humboldt-Universität zu Berlin
meyerhenke@hu-berlin.de

Keywords: algorithm engineering, design and analysis of algorithms, epistemology

1 Abstract

The design, analysis, and evaluation of algorithms have been a major concern of *algorithm engineering* ([1]), which embraces both theoretical and empirical contributions. An *algorithm* is a well-defined finite sequence of computational steps that transforms some input into some output [2]. Today's research and development of algorithms is not restricted to the field of algorithm engineering, but occurs in many 'data science'-related disciplines, such as Machine Learning, Operational Research, Metaheuristics and many others. For 2022 only, Google Scholar lists more than 72,000 research papers with the term *algorithm* in the title[1], highlighting the importance and the richness of research on algorithms.

As a first challenge, we observe that a review of the practice of algorithm engineering is missing. Much of methodological considerations seem to be implicitly clear within the domains of major conferences such as the European Conference on Operational Research (EURO), Very Large Databases (VLDB), Neural Information Processing Systems (NeurIPS), European Symposium on Algorithms (ESA), or Visualization (VIS). On the one hand, it might be argued that such implicit standards reflect a consensus in their

[1] https://scholar.google.com/scholar?as_q=&as_epq=algorithm&as_oq=&as_eq=&as_occt=title&as_sauthors=&as_publication=&as_ylo=2022&as_yhi=2022.

respective fields, which can be interpreted as a sign of maturity. Kuhn [3] calls this *normal science* to signify the incremental work in an established, consensual paradigm. On the other hand, the requirements, e.g., for evaluating algorithmic contributions can differ quite considerably across algorithm-oriented disciplines. While there are good reasons for these differences, the question is whether a common scientific framework can integrate these differences in requirements, such that it might help different algorithm-oriented disciplines to transfer methodological advances.

As a second challenge, we observe that many areas of algorithm engineering have turned to competitive testing with the help of benchmark data sets with established quality metrics to tease out the last epsilon of improvement. As a consequence, most of the knowledge on algorithms is of a comparative nature - which algorithm performs best - with limited research producing explanatory knowledge - why is an algorithm performing well [4]. A common scientific framework on algorithm engineering would provide the necessary meta arguments for authors to deviate from common practice and expand algorithmic research towards producing explanatory knowledge.

The main contribution of this study is the development of such a common research framework for algorithm engineering. Our framework builds on three areas discussed in the philosophy of science: ontology, epistemology and methodology. In essence, *ontology* describes algorithm engineering as being concerned with algorithmic problems, algorithmic tasks, algorithm designs and algorithm implementations. *Epistemology* describes the body of knowledge of algorithm engineering as a collection of prescriptive and descriptive knowledge, residing in World 3 of Popper's Three Worlds model. *Methodology* refers to the steps how we can systematically enhance our knowledge of specific algorithms.

Our framework has important methodological implications for researching algorithms. First, it offers a theoretical foundation for algorithm engineering grounded in the philosophy of science. This foundation clarifies how algorithms can be studied as well as the nature and different types of knowledge that can be obtained. Second, our framework provides guidance for the methodological evaluation of newly proposed algorithms. Authors can use considerations that are related to the different validity concerns as a starting point. Third, this methodological grounding offers a framework for conducting research on algorithms in various 'data science'-related domains in a way that supports an incremental research path. The concepts tied together in our framework might eventually help to foster knowledge transfer across disciplines on how to research algorithms.

References

[1] Peter Sanders (2009). Algorithm engineering–an attempt at a definition. *Efficient Algorithms*. Springer, 321–340.

[2] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein (2009). Introduction to algorithms. MIT press.

[3] Thomas S Kuhn (2012). The structure of scientific revolutions. 50th Anniversary Edition, University of Chicago Press.

[4] John N Hooker (1995). Testing heuristics: We have it all wrong. *Journal of heuristics* 1, 33–42.

Hexaly Optimizer 12.5: interval based modeling and perfomance improvements for routing and scheduling problems

J. Darlay
Hexaly, 251 Boulevard Pereire, 75017 Paris, France
jdarlay@hexaly.com

Keywords: Optimization solver, Routing, Scheduling

1 Introduction

Hexaly Optimizer, previously LocalSolver, is a global solver based on exact and heuristic techniques. Its modeling formalism enables it to accept any model using the usual mathematical operators (arithmetic, logical, relational, etc.) with continuous, integer or set variables. The solver uses exact techniques to obtain lower bounds and prove optimality. It also run heuristics in parallel to quickly find good solutions and to handle larger instances.

The talk will present the new features of version 12.5, including a new modeling formalism for scheduling problems and new set operators to better model exclusion constraints in very large models. This new version brings performance improvements on scheduling and vehicle routing problems, which will be illustrated on literature benchmarks.

2 Modeling

The `list` decisions can be used to represent an ordered sequence of integers, such as the list of tasks to be processed on a machine. In order to represent the start and end dates of a task, an `interval` variable has been added to the modeling formalism. The relational operators $<$ and $>$ have been extended to intervals with the semantics $A < B$, if the interval A ends before the interval B starts. The expressions `start(`A`)`, `end(`A`)`, `length(`A`)` and `contains(`A`,` t`)` respectively represent the start date, the end date, the length of the interval A and the expression "$t \in A$". The combination of `list` and `interval` variables enables disjunctive and cumulative scheduling problems to be modeled efficiently.

The `distinct` expression was introduced to return the list of unique elements of a collection. This expression is particularly useful in assignment problems in which a container can only receive one type of product. This can be modeled by requiring that

the cardinality of distinct object types in the container to be less than one: `constraint count(distinct(set, i => type[i])) <= 1`.

The `intersect` expression is used to represent the intersection of two collections. Exclusions of pairs of objects in the same collection can be modeled in a compact way by constraining the cardinality of the intersection between a collection and the collection of prohibited values: `constraint count(intersect(set, {1, 2, 5})) <= 1`.

The `sort` expression is used to sort a collection according to a criterion. It is useful in stochastic optimization when uncertainty is modeled by scenarios and the user wishes to minimize the value of a particular quantile.

3 Primal heuristics

Set based modeling is particularly adapted to vehicle routing problems. A `list` variable naturally modeling the sequence of customers to visit for a truck. Hexaly Optimizer 12.5 brings performance improvements with the addition of new *destroy & repair* type neighborhoods on *pickup & delivery* problems with time windows (PDPTW). For these problems, the average gap on classical instances in the literature falls from 8% to 2.5% in 60s. A second neighborhood based on the ejection chain principle [1] extends the performance obtained to PDPTW variants with interdiction constraints between vehicles and trucks.

The `interval` variables have been introduced to simplify the writing of scheduling problems while maintaining the performance of previous versions (less than 3% gap in 60s for instances with several hundred tasks: jobshop, flexible jobshop, RCPSP, etc.). Solving techniques have been extended to handle pseudo-preemptive scheduling as well as producer-consumer scheduling.

4 Lower bounds

PDPTW problems and heterogeneous variants of vehicle routing problems are now automatically reformulated into compact integer linear programs [2] for solving to optimality small instances. This compact reformulation is solved in parallel with the cut generation reformulations already present in the solver for computing lower bounds.

Bounds for scheduling problems have been reinforced by the exact resolution of combinatorial relaxations. The idea here is to automatically simplify scheduling problems by eliminating constraints. The resulting problem is solved to optimality by efficient dedicated algorithms. This improves the bounds computed on all scheduling problems involving the minimization of a weighted sum of end dates.

References

[1] F. Gardi, T. Benoist, J. Darlay, B. Estellon, et R. Megel. *Mathematical Programming Solver Based on Local Search*, Wiley, 2014

[2] H. D. Sherali, P. J. Driscoll On Tightening the Relaxations of Miller-Tucker-Zemlin Formulations for Asymmetric Traveling Salesman Problems. *Operations Research 2002 50(4):656-669.*

[3] P. Laborie. Bornes rapides pour l'ordonnancement dans LocalSolver. *ROADEF 2023.*

Bandit Algorithms for Neighborhood Selection in Local Search

S. Michelini, R. De Landtsheer
Combinatorial Algorithmics, CETIC research center, Charleroi
stefano.michelini@cetic.be, renaud.delandtsheer@cetic.be

M. Silva-Muñoz, A. Franzin
Université Libre de Bruxelles (ULB)
moises.silva.munoz@ulb.be, alberto.franzin@ulb.be

A. Delecluse
TRAIL, ICTM, UCLouvain
augustin.delecluse@uclouvain.be

Keywords: Local search, Bandit algorithms, Reinforcement learning

1 Neighborhood selection

When tackling a combinatorial optimization (CO) problem using local search, choosing the right neighborhood functions is crucial for good performance. This is especially true for methods such as variable neighborhood search, where different neighborhood functions are used sequentially. Common selection policies include *round robin*, that iterates through all the available neighbourhoods, and *hill climbing*, that prioritizes the ones that allow to obtain the best improvement. The *best slope first* policy prioritizes the neighborhoods that exhibit the best improvement in solution quality relative to the elapsed computing time. The latter is one of the default policies in OscaR.cbls, the local search framework for CO used in this study [1].

2 Bandit algorithms for neighbourhood selection

In recent years, the application of machine learning techniques for the acceleration of more traditional techniques to solve CO problems has garnered considerable attention [2]. In particular, reinforcement learning (RL) models the behaviour of an agent interacting with an *environment*. The agent's actions change the state of the environment, and are each associated to a *reward*. The goal is to learn a policy of actions that maximizes the total reward.

One class of such methods is multi-armed bandit algorithms, where each action is a choice among a set of competing alternatives. The agent needs to learn the optimal schedule of a limited amount of actions. Recent examples of hybdrization between CO and bandit algorithms include [3], where an agent dynamically selects the variable to branch on during the exploration of a search tree, or [4], where the choice is among several heuristics for branch-and-bound exploration.

Inspired by [4], we define the environment as the state of the local search at each iteration and the action space as the set of available neighbourhoods. The optimal policy is the schedule of neighborhoods that yields the best possible solution. The policy is learned *online*, that is, during the search, without a preliminary training phase. We define several reward schemes, such as rewards depending on the presence of improvement, or proportional to the improvement in solution quality, with different frequencies of update. We also implement different bandit algorithms, including the popular ϵ-greedy and UCB, as well as custom implementations. We evaluate whether this approach can define a neighborhood schedule that outperforms the default best slope first method.

3 Experiments

We evaluated these algorithms on the pickup and delivery problem with time windows, for which several neighborhoods are defined. In our experiments we used the benchmark set by Li and Lim [5].

Preliminary results show that some of these algorithms, notably UCB and ϵ-greedy, perform at least as well as the best slope first method, while other alternatives exhibit higher variance in solution quality. These results may anyway be improved by exploiting the higher flexibility of these methods, entailed by their parameterization, with respect to the best slope first policy.

Further studies are therefore aimed at enhancing overall performance by following several directions, such as introducing a preliminary learning phase to optimize algorithmic parameters, as well as refining the definition of the reward schemes. The inclusion of different bandit algorithms is also under consideration. Additionally, since the implementation of these methods is modular and readily available for other problems defined in OscaR.cbls, further research may include the evaluation of their performance on different classes of CO problems.

Acknowledgement This work is supported by the Service Public de Wallonie Recherche under grant n°2010235 - ARIAC. The authors are active members of the TRAIL Institute (Trusted AI Labs, Fédération Wallonie-Bruxelles, Belgium).

References

[1] De Landtsheer, R., and Ponsard, C. (2013) *OscaR.cbls: an open source framework for constraint-based local search.* ORBEL 2013.

[2] Bengio, Y., Lodi, A., and Prouvost, A. (2021). *Machine learning for combinatorial optimization: A methodological tour d'horizon.* EJOR.

[3] Chalumeau F. *et al.* (2021) *Seapearl: A constraint programming solver guided by reinforcement learning*. CPAIOR 2021.

[4] Chmiela, A., Gleixner, A., Lichocki, P., Pokutta, S. (2023). *Online Learning for Scheduling MIP Heuristics*. CPAIOR 2023.

[5] Li, H., and Lim, A. (2001). *A metaheuristic for the pickup and delivery problem with time windows*. ICTAI 2001.

50 years of metaheuristics

K. Sörensen
University of Antwerp, Belgium
kenneth.sorensen@uantwerpen.be

R. Martí
Universidad de Valencia, Spain
rafael.marti@uv.es

M. Sevaux
Université Bretagne Sud, France
marc.sevaux@univ-ubs.fr

Keywords: Metaheuristics, history

In this talk, we examine the significant advancements in heuristic methods for optimization over the past 50 years. Our focus is primarily on the European perspective and we undertake a critical analysis of the key findings and contributions in this field. We begin by exploring the origins of heuristics, which can be traced back to classical philosophers, and then proceed to trace the historical development of heuristics and metaheuristics in the field of operations research. We highlight the major milestones along this path, including the latest efforts to combine metaheuristics with machine learning. A particular emphasis is placed on the theories that have revolutionized our understanding of problem solving, as well as the role played by the European Journal of Operational Research in advancing these theories. Our approach centers on methodologies and their connections to related areas, enabling us to identify potential avenues for future research.

Multi-Task Learning: a method for grouping tasks without training a network.

Y. Morello, S. Verboven
VUB, Data Analytics Lab
yoann.morello@vub.be, sam.verboven@vub.be

Keywords: Machine Learning, Multi-Task Learning

1 Background

Multi-Task Learning (MTL) is a machine learning approach where a single model is designed to perform multiple tasks simultaneously. It aims at leveraging the interdependencies between different tasks to learn more efficiently. A significant challenge in MTL is the strategic grouping of tasks. Poorly chosen groups of tasks fail to achieve any enhancement. Somewhat related, yet distinct, Hu et al. [2022] has articulated that MTL may be inherently more prone than Single-Task Learning (STL) to the acquisition of spurious correlations, especially in certain configurations of task pairings. Spurious correlations are those which are not causal, as for instance in the case of unbalanced classes.

2 Prior Work

To date, the extraction of relevant information from datasets for task grouping stands as a pivotal process. Even in the context of low computational cost methods, it necessitates the complete training of a neural network. Additionally, the issue of spurious correlations introduces a significant challenge. To the best of our knowledge, this factor has consistently been overlooked in the methodologies adopted for task grouping in MTL, pointing to a gap in the current understanding and application of these systems.

3 Our work

In our work, we propose a method for grouping tasks which does not require preliminary training of a neural network. Furthermore it provides some insights into the impact of potentially spurious correlations. This is achieved by introducing a novel family of measures, with a focus on one particular measure, designed to explore databases and

their subsets. This exploration aims to compare the learning dynamics of neural networks trained on various subsets. Our methodology is inspired by two distinct research trends. Firstly, the work of Kvinge et al. [2021] provides a versatile set of measures for evaluating models in relation to data, leveraging some tools from mathematical category theory. Secondly, the investigations of Saxe et al. [2019] into the differential equations of back-propagation in linear architectures reveal that their learning dynamics, under some strong assumptions, depends only on the Singular Value Decomposition (SVD) of some covariance matrices. These findings hold empirical significance even for more complex, non-linear network architectures. Building on these, our measure introduces a method to quantify the impact of removing a task from a group of tasks or a subset of data points from a set. Specifically, it assesses the effect on the Singular Value Decomposition compression of various correlation matrices. Contrasting with existing methods, our approach yields comparable clustering recommendations through experimental replication, with the notable advantage of eliminating the need for network training. Additionally, our method demonstrates, both theoretically and empirically, the ability to discern effects of features from those arising due to spurious correlations. This includes imbalances in class correlations that persist even in ”balanced” datasets.

4 Its implications

Our approach not only establishes itself as a low computational cost method for task grouping in neural networks but also offers critical insights into the mechanisms of information propagation within these architectures. Significantly, it extends beyond the traditional confines of MTL, contributing to narrowing the gap between data characteristics and neural network learning dynamics. In particular, our findings could provide empirical support for the provocative hypothesis presented in Paul et al. [2021], which suggests that it might be possible to determine the subset of data for pruning directly from the dataset itself, without initial model training.

References

Ziniu Hu, Zhe Zhao, Xinyang Yi, Tiansheng Yao, Lichan Hong, Yizhou Sun, and Ed Chi. Improving multi-task generalization via regularizing spurious correlation. *Advances in Neural Information Processing Systems*, 35:11450–11466, 2022.

Henry Kvinge, Brett Jefferson, Cliff Joslyn, and Emilie Purvine. Sheaves as a framework for understanding and interpreting model fit. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4222–4230, 2021.

Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems*, 34:20596–20607, 2021.

Andrew M Saxe, James L McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019.

Selecting the Optimal Lawful ML Model

M. Hanson
Vrije Universiteit Brussel, Data Analytics Laboratory
Mathias.Hanson@vub.be

S. Verboven
Vrije Universiteit Brussel, Data Analytics Laboratory
Sam.Verboven@vub.be

Keywords: Regulatory Compliance, Machine learning, European Law

1 Introduction

Machine learning (ML) systems developed for and deployed in real-world applications must comply with the applicable legal framework. Ensuring such compliance can be complex. On the one hand, legal experts must identify which laws apply and then determine the correct application thereof. Both tasks are inherently ambiguous as these are questions of interpretation. On the other hand, data scientists must develop technical solutions that integrate legal constraints into ML systems.

In addition to the individual complexity of these two aspects, the design of lawful ML systems raises another key question: how should one identify the optimal lawful model? The literature provides limited guidance on this question. While there is sparse literature that explores technical translations of specific legal principles, see e.g., [1], this literature does not address any approach for identifying the 'optimal lawful model'. In [1], for instance, the authors first translate the EU General Data Protection Regulation's (GDPR) data minimisation requirement [2] to fixed legal constraints and then analyse model performances within those fixed legal constraints without addressing model selection. On the basis of such literature, one could then draw the impression that identifying the optimal lawful model is merely a matter of selecting the best performing model within fixed legal constraints.

However, such an approach fails to recognise that the legal constraints themselves will often be uncertain. Therefore, we propose an integrated approach to enable exploration within flexible legal guardrails, which consists of 3 main elements: (1) legal experts identify multiple potential interpretations and translate those into alternative sets of legal constraints (rather than a single fixed set), each with their individual legal risk; (2) data scientists then construct various models within each plausible set of legal constraints; (3) finally, the optimal model is chosen by balancing both legal risk and predictive quality

through a joint evaluation and selection process. In what follows, a simple example is developed to illustrate our approach.

2 The Case for an Integrated Approach: an illustrative Example

We consider a current legal issue characterised by considerable uncertainty. In this case, we focus on the question: can an ML model be considered 'personal data' under the GDPR and, if so, what would some consequences thereof be?

Some scholars argue a machine learning model can effectively be considered personal data under the GDPR if a the model can be successfully 'attacked' to infer information about natural persons that were included in the training set [3]. However, other scholars argue the opposite view that a machine learning model should not be considered 'personal data' under the GDPR [4].

Depending on the interpretation adopted, there are different legal consequences and thus different legal constraints. If one follows the interpretation that ML models are not 'personal data' under the GDPR, no additional legal constraints are introduced. If one takes the alternative interpretation, additional legal constraints will be incorporated. To further develop the example, we briefly set out some potential constraints as identified in the article [3]:

- **Constraint 1:** the capability to remove specific personal data from a ML model upon request [5];
- **Constraint 2:** resiliency to attacks aimed at inverting the model; and
- **Constraint 3:** the capability to automatically remove personal data that is no longer necessary to retain.

In other words, the interpretation of ML models as personal data may introduce significant constraints that could limit model performance if imposed too rigidly at the outset. Rather than locking in stringent restrictions initially, the integrated approach involves legal experts evaluating, for each of the constraints, whether there exist other legally valid approaches. This leads to various constraint sets, each with associated legal risk levels (e.g., one set with constraints 1 and 2, another with all three constraints, etc.). Data scientists then construct multiple models within each specific set of constraints (e.g., Models A and B in set 1, Models C and D in set 2, etc.). Finally, legal and technical experts jointly evaluate the models by weighing the legal risk profiles and their performance to select the optimal lawful model (e.g., Model C may be selected over A, B and D due to significantly lower legal risk than A and B and higher performance than D).

3 Conclusion

We argue that the decision to select the optimal lawful ML model should recognise that the legal constraints are in themselves uncertain and, thus, be based on a trade-off between legal risk and model performance. This contrasts with the approach that model selection should solely be a matter of comparing model performance within a pre-defined

set of legal constraints. In future work, the aim is to empirically show potential benefits of our proposed approach.

References

[1] Biega, A.J., Potash, P., Daumé, H., Diaz, F. & Michèle Finck (2020). Operationalizing the Legal Principle of Data Minimization for Personalization. In Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, 399–408. https://doi.org/10.1145/3397271.3401034

[2] Consolidated text: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016. Article 5(1)(c).

[3] Veale, M., Binns, R. & Edwards L. (2018). Algorithms that remember: model inversion attacks and data protection law. Phil. Trans. R. Soc. A, 376(2133), Article 20180083. http://dx.doi.org/10.1098/rsta.2018.0083

[4] Leiser, M. R. & Dechesne, F. (2020). Governing machine-learning models: challenging the personal data presumption. International Data Privacy Law, 10(3), 187–200. https://doi.org/10.1093/idpl/ipaa009

[5] Xu, H., Zhu, T., Zhang, L., Zhou, W. & Yu, P.S. (2023). Machine Unlearning: A Survey. ACM Computing Surveys, 56(1), Article 9, 1–36. https://doi.org/10.1145/3603620

Personalizing Discounts Using Learning to Rank

T. Vanderschueren
KU Leuven, University of Antwerp

F. Moraes
Booking.com

H. Proença
Booking.com

Keywords: Causal Inference, Uplift Modeling, Learning to Rank

Efficiently allocating treatments with a budget constraint constitutes an important challenge across various domains. In marketing, for example, the use of promotions to target potential customers and boost conversions is limited by the available budget. While much research focuses on estimating causal effects, there is relatively limited work on learning to allocate treatments while considering the operational context. Existing methods for uplift modeling or causal inference primarily estimate treatment effects, without considering how this relates to a profit maximizing allocation policy that respects budget constraints. The potential downside of using these methods is that the resulting predictive model is not aligned with the operational context. Therefore, prediction errors are propagated to the optimization of the budget allocation problem, subsequently leading to a suboptimal allocation policy. We propose an alternative approach based on learning to rank. Our proposed methodology directly learns an allocation policy by prioritizing instances in terms of their incremental profit. We propose an efficient sampling procedure for the optimization of the ranking model to scale our methodology to large-scale data sets. Theoretically, we show how learning to rank can maximize the area under a policy's incremental profit curve. Empirically, we validate our methodology and show its effectiveness in practice through a series of experiments on both synthetic and real-world data.

Discerning customer journey types through dynamic modeling of usage Patterns

J. Sanchez Ramirez
IESEG School of Management, Univ. Lille, Ghent University,
j.sanchezramirez@ieseg.fr

K. Coussement, Arno De Caigny
IESEG School of Management,
k.coussement@ieseg.fr, a.de-caigny@ieseg.fr

D. F. Benoit
Ghent University,
dries.benoit@ugent.be

L. Waardenburg
ESSEC Business School,
waardenburg@essec.edu

Keywords: Usage patterns, customer journey, graph-based representation

With the evolution of big data and customer analytics, businesses are storing a diverse set of data sources resulting from multiple touchpoints, including marketing interactions, support, training, transaction history, and usage. Integrating this information into a holistic view of the customer journey provides insights into consumer decisions, allowing analysis of short and long-term impacts of marketing interactions and guiding program designs for the entire customer lifecycle.

In this project involving real-life use case, we first explore the integration of different data sources as part of the customer journey (CJ). Second, we investigate the integration of usage data through the CJ to identify different usage patterns. Finally, we evaluate the impact of these usage patterns on business-relevant metrics such as customer lifetime value (CLV) and churn rate.

This research contributes to the literature on the exploration of usage data in a network-based setting. It aims to capture relevant aspects such as the temporal structure of usage and the integration of heterogeneous inputs. Our research has the potential to provide a transferable methodology capable of capturing and incorporating usage patterns in different applications of business analytics models, enhancing their adaptability and effectiveness across diverse contexts.

Public Debt Dynamics, Debt Sustainability and the Role of Debt Relief Initiatives: Empirical from Ethiopia, DSSI & CF

T. Merra
PhD Candidate in Applied Economics, University of Antwerp
tesfaye.merra@student.uantwerpen.be

D. Cassimon
University of Antwerp, Chair of IOB

J. Abafita
Associate Professor, Jimma University of Department of Economics

Keywords: External debt, Debt servicing, Debt service suspension, Real option

This paper aims to explore the main drivers of the renewed build-up of public debt and the role of debt service suspension initiative (DSSI) and the common framework (CF) beyond the DSSI in Ethiopia using time series data from 1980 to 2022, paying a more attention to the last decade. Autoregressive distributive lag (ARDL) bound test and error correction model was employed in dealing with the drivers of debt accumulation, and followed by a real option approach to assess the role of DSSI and CF. Using a real option method, we assessed the decision to opt participation (benefits of participation) or not to participate (fearing costs of participation) to these debt relief initiatives. The result of the co-integration test of the ARDL bound test reveals the existence of multiple long-run equilibrium relationship between the debt variables and debt accumulation. In the long run the main drivers of debt accumulation were debt servicing, inflation and national savings where as in the short run debt servicing, exchange rate, inflation and primary budget deficit remains the main determinant of debt build ups. The empirical result shows that external debt servicing has a significant and positive long run relationship with external debt accumulation. The elasticity of external debt stock with respect to debt servicing is 0.8393, that implies, a percentage change in debt servicing will increase the external debt accumulation by 84 percent and the P-value (p-0.000), reveals that debt servicing is statically significant at 1 percent level in affecting debt accumulation in the country. As a result, the positive effect of current debt service on debt accumulation was as expected, based on both theoretical and empirical basis. The result also depicts that inflation and saving were statistically significant and have long run relationship that affect debt build

Bankruptcy Problems over Time

M. Vollebergh
Vrije Universiteit Amsterdam, Department of Operations Analytics
m.a.j.vollebergh@vu.nl

W. Dullaert, A. Estévez-Fernández & Y. Ghiami
Vrije Universiteit Amsterdam, Department of Operations Analytics

Keywords: game theory, bankruptcy, allocation problems

Allocation problems arise often in practice: ranging from profit allocation in an intermodal transport system (Algaba et al., 2019), to sharing water in a transboundary river (Ansink and Weikard, 2012) and the allocation of vaccines over multiple populations (Duijzer et al., 2018). These problems can be modeled in the framework of bankruptcy problems. A bankruptcy problem arises when a group of agents each claims a portion of an estate that cannot satisfy the sum of all claims, leading to the question of how to fairly divide the available estate among the agents. A solution to bankruptcy problems is given by so-called bankruptcy rules, which specify the part of the estate that each agent should get. In the traditional bankruptcy problem, there is a single decision made at a single point in time. Looking at the above applications, one can see that these are actually problems where allocations have to be made over time.

In this presentation, we introduce the concept of bankruptcy over time in the Multi-period Bankruptcy Problem. We also present three different perspectives on bankruptcy over time. For each perspective, we generalize existing bankruptcy rules, examine and prove their properties, and derive efficient algorithms to apply them.

To the best of our knowledge, bankruptcy rules over time have not been considered in the literature. In Chen and Thomas (2018), the time aspect is considered for what they call common allocation rules, which are the same as the common bankruptcy rules. However, Chen and Thomas (2018) do not take into account the loss that each agent has taken in previous periods in the proposed common allocation rules. Recently, sequential bankruptcy problems have been introduced in Estévez-Fernández et al. (2021). This generalization of the Bankruptcy Problem covers, among others, the problems of sharing water in a transboundary river. Although in this paper the estate is divided into multiple parts, it still considers a single decision made at some point in time. Our approach therefore goes further than this generalization by considering multiple decisions that will be taken successively.

The perspectives on bankruptcy over time that we present each consider the claims in each time period differently. The way the estate is considered, however, remains the same. In each time period, the available estate is fully utilized.

The first perspective on bankruptcy over time considers each time period separately: only the claims given for the current period are taken into account for determining the allocations. The second perspective also takes the difference between the agents' claims and the agents' awards in the previous periods into consideration, while still making a distinction between this 'memory' and the claims of the current period. The third perspective considers for each agent the difference between that agent's claims and that agent's awards in the previous periods and the claim given for the current period together as one claim.

For each of these perspectives, we present mechanisms to generalize existing bankruptcy rules to the multi-period setting. Moreover, we formulate the properties of these multi-period bankruptcy rules. This is done in such a way that the properties are inherited: if a bankruptcy rule satisfies this property, then the multi-period generalization of this rule also satisfies the property. Finally, we present efficient algorithms to compute the results of a bankruptcy rule by using recursive methods and giving explicit formulas to calculate the parameters in these rules.

By introducing the Multi-period Bankruptcy Problem, presenting and computing multi-period bankruptcy rules, we hope to extend the usefulness of bank-ruptcy rules in practice and open avenues for further research on multi-period bankruptcy problems.

References

[1] Algaba, E., Fragnelli, V., Llorca, N., & Sánchez-Soriano, J. (2019). Horizontal cooperation in a multimodal public transport system: The profit allocation problem. European Journal of Operational Research, 275 (2), 659–665.

[2] Ansink, E., & Weikard, H.-P. (2012). Sequential sharing rules for river sharing problems. Social Choice and Welfare, 38 (2), 187–210.

[3] Chen, C.M., & Thomas, D. J. (2018). Inventory allocation in the presence of service-level agreements. Production and Operations Management, 27 (3), 553–577.

[4] Duijzer, L. E., van Jaarsveld, W. L., Wallinga, J., & Dekker, R. (2018). Dose-optimal vaccine allocation over multiple populations. Production and Operations Management, 27 (1), 143–159.

[5] Estévez-Fernández, A., Giménez-Gómez, J.-M., & Solís-Baltodano, M. J. (2021). Sequential bankruptcy problems. European Journal of Operational Research, 292 (1), 388–395.

Rawlsian Assignments

T. Demeulemeester
KU Leuven, Research Center for Operations Research & Statistics
Ghent University, Department of Business Informatics and Operations Management
tom.demeulemeester@kuleuven.be

J. S. Pereyra
Universidad de Montevideo, Department of Economics and Business
jspereyra@um.edu.uy

Keywords: Random assignment, Fairness, Rawls

1 Abstract

We study the assignment of indivisible goods to individuals when monetary transfers are not allowed. Previous literature has mainly focused on efficiency (from both ex-ante and ex-post perspectives) and *individually* fair assignments. As a result, egalitarian concerns have been overlooked. We draw inspiration from the assignment of apartments in housing cooperatives, where families consider egalitarianism of assignments as a first-order requirement. Specifically, they aim to avoid situations where some families receive their most preferred apartments while others are assigned options ranked very low in their preferences.

Building on Rawls' concept of fairness, we introduce the notion of a Rawlsian assignment. We prove that a unique Rawlsian assignment always exists. Furthermore, the Rawlsian rule is efficient and anonymous. To compute the Rawlsian assignment, we propose a polynomial-time algorithm based on linear programming.

To illustrate our analysis, we use preference data from 24 housing cooperatives from Uruguay. We show that the Rawlsian rule substantially improves, from an egalitarian perspective, both the well-studied probabilistic serial rule (Bogomolnaia & Moulin, 2001), and the rule currently used to assign apartments in the housing cooperatives.

2 References

[1] Bogomolnaia, A., & Moulin, H. (2001). A new solution to the random assignment problem. Journal of Economic theory, 100(2), 295-328.

Algorithmic Solutions for Maximizing Shareable Costs

R. Zou
Northwestern Polytechnical University, School of Mathematics and Statistics
rongzou20@gmail.com

B. Lin
KU Leuven, ORSTAT, Faculty of Economics and Business (FEB)
boyue.lin@kuleuven.be

M. Uetz
University of Twente, Department of Applied Mathematics
m.uetz@utwente.nl

M. Walter
University of Twente, Department of Applied Mathematics
m.walter@utwente.nl

Keywords: Linear Programming, Computational Complexity, Cost Sharing, Minimum Spanning Tree Game

We addresses the linear optimization problem to maximize the total costs that can be shared among a group of agents, while maintaining stability in the sense of the core constraints of a cooperative transferable utility game, or TU game. When maximizing total shareable costs, the cost shares must satisfy all constraints that define the core of a TU game, except for being budget balanced.

We first study the computational complexity of this optimization problem, its relation to optimization over the core itself, and its equivalence to other, minimal core relaxations that have been proposed earlier.

Then we address minimum cost spanning tree (MST) games as an example for a class of cost sharing games with non-empty core. While submodular cost functions yield efficient algorithms to maximize shareable costs, MST games have cost functions that are subadditive, but generally not submodular. Nevertheless, it is well known that cost shares in the core of MST games can be found efficiently. In contrast, we show that the maximization of shareable costs is NP-hard for MST games and derive a 2-approximation algorithm. Our work opens several directions for future research.

Dynamic Assignment for a Crowdshipping Platform

E. Innocente
Université catholique de Louvain, Center for Operations Research and Econometrics (CORE)
emma.innocente@uclouvain.be

J.-S. Tancrez
Université catholique de Louvain, Center for Operations Research and Econometrics (CORE)
js.tancrez@uclouvain.be

Keywords: Crowdshipping, Dynamic Assignment, Approximate Dynamic Programming

Driven by the growth of online shopping, parcel delivery has grown rapidly in recent years, resulting in an increased demand for cost-efficient last-mile delivery [3]. At the same time, the question of sustainability is increasingly present in logistics and in society in general.

Crowdshipping applies the concept of crowdsourcing to the personalized delivery of freight, turning ordinary citizens into couriers for the distribution of small items. In this collaborative delivery system, individuals already traveling from an origin to a destination take charge of all or part of the delivery, taking a package along with them and making a stop along the way to drop it off. In a successful implementation, crowdshipping reduces the number of freight delivery trucks, benefiting companies by reducing their delivery costs, as well as improving sustainability [4].

The matching of crowdshippers to parcels as they appear in real time is key for crowdshipping to yield its intended effects. Indeed, the high degree of uncertainty inherent to crowdshipping, primarily in crowdshippers' and parcels' availability, can result in high operational costs and low service reliability [2]. In determining the assignment of crowdshippers to parcels over time, the main questions are as follows: (1) To which crowdshipper should we assign the parcels? (2) Should we postpone some delivery in the hope for a future crowdshipper with a shorter detour? (3) How can we consider the future uncertainty and downstream cost in real-time decision making? (4) Can we increase the total number of served parcels and/or decrease delivery distance by making smarter assignment decisions?

Interestingly, the majority of research on crowdsourced transportation does not directly address the challenges of uncertain availability, but rather assume perfect information. Studies that incorporate the uncertainty in crowdshipping operations are scarce, and only

few of those consider a dynamic problem setting (e.g., [1]; [5]; [7]; and [2]). Recent works by Mousavi et al. [2] and Dayarian and Savelsbergh [1] are the closest to ours. However, both consider in-store customers as crowdshippers to deliver groceries within few hours.

Our article focuses on crowdshippers with any origin and destination, and longer distances. Crowdshippers are known in advance, enabling planning and decision-making more ahead of time. Exogenous information relative to the arrival and disappearance of parcels and crowdshippers is observed throughout the day. At the beginning of each day, decisions regarding the assignment of crowdshippers to parcels are made among those available. Then at the end of the day, based on the decisions and observed exogenous information, the state of the system is updated. The dynamic assignment problem considered in our article therefore consists of solving a sequence of crowdshipper-to-parcel assignment problems over time, with both crowdshippers and parcels arriving randomly over time. By their very nature, crowdshippers will only be available a single day, while parcels have a delivery window. This problem can be classified as a discrete time sequential stochastic decision problem, that we formulate as a discrete time Markov decision process.

Due to the "curse of dimensionality", the model is not solved with conventional dynamic programming methods. The proposed algorithm uses an Approximate Dynamic Programming (ADP) approach based on Value Function Approximation (VFA) [6]. It is based on offline learning of value functions through simulation of the problem's Markov decision process. This adaptive learning algorithm provides nonmyopic behavior yet requires only solving sequences of assignment problems no larger than would be required with a myopic algorithm.

Our first numerical results show the efficiency of our algorithm in comparison with myopic solutions on small instances. We compare the performance of our algorithm against both myopic and perfect-information approaches. The comparison is made in terms of the number of parcels handled by crowdshipping, as well as of the reduction of the delivery distance.

References

[1] Dayarian, I. and Savelsbergh, M. (2020). Crowdshipping and Same-day Delivery: Employing In-store Customers to Deliver Online Orders. *Production and Operations Management*, 29(9):2153–2174.

[2] Mousavi, K., Bodur, M., Cevik, M., and Roorda, M. J. (2021). Approximate Dynamic Programming for Crowd-shipping with In-store Customers.

[3] Perez, S. (2020). COVID-19 pandemic accelerated shift to e commerce by 5 years, new report says. Technical report.

[4] Punel, A. and Stathopoulos, A. (2017). Modeling the acceptability of crowdsourced goods deliveries: Role of context and experience effects. *Transportation Research Part E: Logistics and Transportation Review*, 105:18–38.

[5] Raviv, T. and Tenzer, E. Z. (2018). Crowd-shipping of small parcels in a physical internet.

[6] Spivey, M. Z. and Powell, W. B. (2004). The Dynamic Assignment Problem. *Transportation Science*, 38(4):399–419.

[7] Yıldız, B. (2021). Express package routing problem with occasional couriers. *Transportation Research Part C: Emerging Technologies*, 123:102994.

Dynamic Pricing, Admission and Repositioning in Ride-Sharing Networks

T. De Munck
UCLouvain, Center for Operations Research and Econometrics
thomas.demunck@uclouvain.be

P. Chevalier
UCLouvain, Center for Operations Research and Econometrics
philippe.chevalier@uclouvain.be

J.-S. Tancrez
UCLouvain, Center for Operations Research and Econometrics
jean-sebastien.tancrez@uclouvain.be

Keywords: Ride-sharing platforms, deep reinforcement learning, network management.

1 Problem Description

Over recent years, ride-sharing platforms (e.g., Uber, Lyft, or Bolt) have considered various levers to balance supply and demand in their network. This work explores three categories of decisions made on ride-sharing platforms: pricing, customer admission, and driver repositioning decisions. The goal is to evaluate the value of these decisions and understand how they influence the dynamics of ride-sharing networks.

2 Methodology

To this aim, we develop a model formulated as a discrete-time Markov decision process. Customers are impatient and arrive with varying willingness to pay. Drivers can decide to self-relocate to maximize their expected earnings, given the platform's pricing policy. The platform operates over a network of locations characterized by supply-demand imbalances. In this setting, the objective is to find a dynamic policy that maximizes the revenues while minimizing the costs for driver repositioning and customer rejection.

Due to the "curse of dimensionality" (Powell, 2007), the model cannot be solved to optimality with conventional methods such as dynamic programming. We address this issue by deploying a methodology that combines mathematical optimization and deep reinforcement learning (DRL). Specifically, our approach proceeds in three stages to find a

near-optimal policy. First, the stochastic model is formulated as a deterministic optimization problem to obtain a heuristic policy. Second, neural networks are trained to reproduce the heuristic policy. Third, the policy accounts for stochastic components of the environment using deep reinforcement learning.

3 Results and Contributions

In a numerical study, we apply our approach to real-world data from the New York City Taxi and Limousine Commission. We first demonstrate the efficiency of our approach in comparison with regular deep reinforcement learning, both in terms of final performance and computational effort. We then explore several managerial insights by analyzing the decisions taken by the near-optimal policy. Namely, admission and repositioning control are used in different situations and can lead to positive implications for the platform, the customers, and the service providers.

The contribution of this work is twofold. From a modeling perspective, we develop a model that optimizes joint pricing, customer admission, and driver repositioning in an extensive setting. In contrast to existing studies (e.g, Afeche et al., 2023), our framework includes multiple layers of complexity that capture (i) spatial-temporal supply-demand imbalances, (ii) stochastic demand and travel times, and (iii) idle drivers with self-relocating behavior. From a methodological perspective, we propose an original approach that can significantly reduce computational effort. To date, there is still limited research (e.g., Oroojlooyjadid et al., 2022, De Moor et al., 2022) that leverages transfer learning methods in operational models. Our methodology stands out by its recourse to mathematical optimization as a way to fasten the training process of the reinforcement learning algorithm.

4 References

[1] Afèche, P., Liu, Z., & Maglaras, C. (2023). Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Manufacturing & Service Operations Management, 25(5), 1890-1908.

[2] De Moor, B. J., Gijsbrechts, J., & Boute, R. N. (2022). Reward shaping to improve the performance of deep reinforcement learning in perishable inventory management. European Journal of Operational Research, 301(2), 535-545.

[3] Oroojlooyjadid, A., Nazari, M., Snyder, L. V., & Takáč, M. (2022). A deep q-network for the beer game: Deep reinforcement learning for inventory optimization. Manufacturing & Service Operations Management, 24(1), 285-304.

[4] Powell, W. B. (2007). Approximate Dynamic Programming: Solving the curses of dimensionality (Vol. 703). John Wiley & Sons.

Metaheuristic algorithm for the design of one-way electric car sharing systems

H. Taki Eddine
Royal Military Academy, Department of Maths
Taki.Harbi@mil.be

L. Ben
Royal Military Academy, Department of Maths
Ben.Lauwens@mil.be

G. Vanden Berghe
KU Leuven, Department of computer science
greet.vanden.berghe@kuleuven.be

Keywords: Optimization, Metaheuristic, Electric Car sharing

Car sharing systems provide users with short-term access to a variety of shared vehicles, including cars, bikes, scooters, or vans. It is an innovative transport strategy that allows users to use the adequate type of vehicle according to their needs and only when required. Various car-sharing systems have been established, falling into three categories [1]: round trip, where the trips start and end at the same station; one-way trip, allowing users to pick up a car from one station and drop it off at another; and finally, free-floating systems where cars are not based at specific stations.
The configuration of these systems has a major impact on the car sharing operators' revenues and their ability of serving maximum number of customers. This configuration consists of several parameters such as the fleet size or the number of cars, their types, their battery or fuel tank capacity, the number of the stations and their locations in the city.
In this work, we present a metaheuristic method to design an effective one-way electric car sharing system. The aim is to select a set of locations to establish the stations, the initial number of cars at each station and select a feasible trip if exists for each customer. We tested our method on the scenarios published by [2], and our method shows promising results. Next step is to consider the stochasticity in the demand.

Determining areas suitable for demand-responsive transit

M. Wens
KU Leuven, Institute of mobility
maarten.wens@kuleuven.be

P. Vansteenwegen
KU Leuven, Institute of mobility
pieter.vansteenwegen@kuleuven.be

Keywords: Public transport, line planning, population diversity

Many public transit operators notice the limitations of providing a public bus network only consisting of fixed lines. This type of transportation is rigid, inflexible, and assumed to be sub optimal in places with low demand. At the same time, academics have been doing theoretical research on alternative services, such as demand-responsive, on-demand, or paratransit, for decades. An aspect of these flexible services which is currently not often included is where they should be implemented, and how they can synergise with other public transport services [1]. This research tries to fill that gap.

The first step in this research is to design a network of fixed lines based on a a population-based optimisation algorithm. In this algorithm, different priorities are modelled at the same time. As a result, each member of the population provides a different network optimised under different conditions (e.g. different penalisation waiting times). This method allows to determine which parts of the studied area are certainly served well with conventional lines, because these lines return in most of the alternative networks. At the same time, zones where the resulting networks differ in offering a service or not, are interesting candidates to provide an alternative service.

A second benefit from the inclusion of these different objectives inside the optimisation algorithm, is that it can be used to improve traditional population based algorithms. The optimisation of several networks in parallel, each with similar goal but with different weights attributed to different parts of the objective function, leads to a natural diversity in the population. The algorithm still converges, but each subpopulation converges to a slightly different solution.

This algorithm is shown in table 1 to beat the current benchmarks for fixed-line optimisation in decent time. Furthermore, it is capable to optimise the current network of Flemish busses in a case study for ”De Lijn”, the main Flemish bus operator. The case study consist of the Mechelen transport region. In this case study an optimal bus network is found for a network with 700 zones, 500 bus stops, 70 bus lines, 37 train stations and 7 train lines. The cost included are an access cost to the bus stops (walking), a waiting costs which is a function of the frequencies of the lines, an in-vehicle driving cost, and transfer

Synchronized Road Unblocking and Relief Material Distribution for Effective Disaster Response

H. Şahin
Department of Industrial Engineering
TOBB University of Economics and Technology, Ankara, Turkey
halenursahin@gmail.com

A. İrfan Mahmutoğulları
Department of Computer Science, KU Leuven
irfan.mahmutogullari@kuleuven.be

M. Kadan
Department of Industrial Engineering
TOBB University of Economics and Technology, Ankara, Turkey
mervekadan1@gmail.com

Keywords: Humanitarian logistics, Vehicle routing, Relief material distribution, Road unblocking.

Natural disasters, with their unpredictable and large-scale impact, cause significant physical, economic, and social losses. These events generate residual materials known as debris. In disaster areas, debris often blocks roads and impedes transportation. However, prompt access to disaster areas becomes crucial in order to improve the quality of life for affected individuals by efficiently delivering relief materials to those in need. In this study, we address the problem of synchronized unblocking of roads blocked by debris in the disaster areas and the delivery of relief materials to critical locations as quickly as possible. This problem involves coordinating the actions of a single road-clearing machinery and multiple homogeneous relief material delivery vehicles, a scenario not yet realistically addressed in the literature. We propose a mixed-integer linear programming model to solve the problem. Since the proposed model is not able to find a solution in a reasonable time for realistic-sized problems, a heuristic algorithm is developed for large instances. The heuristic is based on an adaptive large-neighborhood framework that produces high-quality solutions. To evaluate the performance of the mixed-integer linear programming model and the heuristic algorithm, a series of computational experiments are conducted using synthetic data that emulate urban street structures. These experiments analyze the impact of roadblocking rates, the number of relief material delivery vehicles, and the number of critical locations on the optimal solution. The findings indicate that the proposed heuristic is effective in terms of both solution time and quality,

Heuristic for an integrated nurse routing and rerostering problem in hospital-at-home

A. Maharani
University of Liège, QuantOM, HEC Liège - Management School
anisha.maharani@uliege.be

V. François
University of Liège, QuantOM, HEC Liège - Management School
veronique.francois@uliege.be

Y. Arda
University of Liège, QuantOM, HEC Liège - Management School
yasemin.arda@uliege.be

Keywords: workforce scheduling and routing problem, nurse rerostering, large neighborhood search

1 Introduction and problem definition

Hospital-at-home (HaH) is an emerging trend that provides care for patients whose medical condition requires admission to a hospital ward, but is sufficiently stable to be treated at home. The following are the expected benefits of such systems: increased capacity of the institutional health care system, potential cost savings, improved patient quality of life, and shorter length of stay [2]. This study focuses on the operational planning decisions in HaH services. Having an acute illness, patients in HaH need intensive care for a limited period. As a result, the patient mix fluctuates. Moreover, resource availability varies over time. Thus, a precise scheduling plan is required to better match the actual resources and the demand.

Given a baseline roster of nurses and the requirements from patients, several operational decisions are taken simultaneously over the planning horizon: select the patients to be admitted, assign nurses to those admitted patients, sequence and schedule the care visits to build the daily tour of each nurse, and decide whether and to what extent the baseline roster should be updated. The studied problem is a combination of task scheduling, nurse routing, and nurse rerostering problems. Generally, the underlying subproblems are solved independently in the literature, as this approach is less complex and thus more computationally practical. However, this sequential decision-making process may lead to

inefficient or even infeasible solutions, since those subproblems are strongly intertwined.

A lexicographic objective function is used to first maximize the number of patients treated at home, and second minimize the total working duration of the nurses. A variety of complex real-world characteristics are considered, yielding a rich integrated problem. The routing subproblem addresses additional constraints associated with home healthcare services, including aligning the medical expertise of nurses with the needs of patients and respecting the time window for treatment administration. When rerostering is necessary, compliance with institutional working hour regulations, encompassing limits on working days and adherence to forbidden shift sequences, is essential. Additionally, the continuity of rostering constraints between consecutive planning horizons is also ensured.

2 Solution method

This work proposes a heuristic solution method to solve the task scheduling, routing, and rerostering subproblems in an integrated manner, aiming to obtain a good HaH operational plan. Infeasible solutions are allowed to be explored during the search by relaxing time windows, rostering, and rerostering constraints. In the first stage of the algorithm, an initial solution is constructed using a greedy insertion mechanism to select which potential patients should be admitted and schedule the care visits of those admitted patients given the baseline roster. The algorithm improves the initial solution in the second phase oscillating among different directions in the search space, i.e. patient-related neighborhoods, route-related neighborhoods, and roster-related neighborhoods. Destroy-and-repair mechanisms based on the large neighborhood search (LNS) approach [1] are employed to increase the number of admitted patients and improve the performance of the routing subproblem. A guided local search is embedded in the heuristic to restore the feasibility of the roster and to refine its performance. The first results obtained by this algorithm will be presented, followed by a discussion about several challenges encountered due to the integration of scheduling, routing, and rerostering decisions.

3 References

[1] Shaw, P. (1998). Using constraint programming and local search methods to solve vehicle routing problems. International conference on principles and practice of constraint programming, 417-431.

[2] Shepperd, S., & Iliffe, S. (2005). Hospital at home versus in-patient hospital care. Cochrane database of systematic reviews, (3).

On the role of model validity, coverage, and overlap in algorithmic fairness Introducing the Minimum Viable Model

C. Bockel-Rickermann
KU Leuven, Faculty of Economics & Business
Christopher.Rickermann@KULeuven.be

T. Vanderschueren
KU Leuven, Faculty of Economics & Business

S. Goethals
University of Antwerp, Faculty of Business & Economics

T. Verdonck
University of Antwerp – imec, Department of Mathematics
KU Leuven, Department of Mathematics

W. Verbeke
KU Leuven, Faculty of Economics & Business

Keywords: Fairness in ML, Causal Inference, Overlap

The increasing abundance of machine learning (ML) applications in business, research, policy-making, medicine, and beyond has led to a growing interest in measuring and protecting algorithmic fairness.

The established literature on fairness in ML has thus far focused on biases in training data which impact the fairness of a model on the output level. Tools to measure fairness include assessments of accuracy and prediction heterogeneity across (sub-)populations of interest. This is grounded in the belief that biases in models stem primarily from biases in the training data, but that the model specification is correct. Such an assumption, however, might not hold in a variety of settings and might worsen analytical models, for example, through enforcing a fairness-accuracy trade-off when it is not needed.

Hence, in this work, we will link fairness in ML to causal inference. The literature on causal inference has discussed a plethora of challenges that hinder the accurate specification of ML models, especially when dealing with observational data. This is a critical threat to algorithmic fairness. While finding a truly causal model might be out of reach in many situations, many concepts from causal inference have the potential to improve

fairness nevertheless. In this regard, we will focus on the concepts of "overlap'' and "coverage''.

Overlap, in the context of treatment effect estimation, and coverage more broadly, refers to the necessity to observe sufficiently similar observations in the training data to make predictions on an unseen observation. If adequate coverage is not assured, the best a model can do is to inter- or extrapolate based on some known concepts of data. This is troublesome whenever new observations might behave differently from known training data. In the process, models might give a confident estimation of an output, yet fail to capture reality. The consequence can be biases in models that do not stem from biases in the generating process of training data, but from assuming that the training data is sufficiently expressive when it is not. For an extreme example, it is false to assume that a model trained on the behavior of dogs, can predict the behavior of cats.

To achieve higher degrees of fairness in ML, we introduce the minimum viable model (MVM). An MVM is a model that has sufficient coverage in the training data for any prediction that it makes. In reverse, an MVM is constrained to only making such predictions where coverage is met. These constraints prevent unreasonable inter- and extrapolations, hinting to practitioners where an ML model should not be used and where it can safely be applied. Therefore, an MVM can ensure better model specification. Rather than adjusting a model globally, we take a fine-grained approach that is observation-specific, and that can be a step before discretionarily and globally adjusting models on the output level.

The concept of an MVM is relevant in a wide range of ML applications. Alongside practical examples, our work presents a definition of an MVM, how it can be built from data, and how it can be used in both prediction- and decision-making contexts. We illustrate that understanding coverage has positive implications for both reaching higher degrees of fairness and improved business operations.

A Pre-processing Approach to Strategic High-Value Customer Representation in Churn Prediction

S. Rahman
Ghent University, Faculty of Economics and Business Administration
shimanto.rahman@ugent.be

B. Janssens
Ghent University, Faculty of Economics and Business Administration
bram.janssens@ugent.be

M. Bogaert
Ghent University, Faculty of Economics and Business Administration
matthias.bogaert@ugent.be

Keywords: Churn, Profit-Driven Analytics, Preprocessing

1 Extended Abstract

In the dynamic landscape of Customer Relationship Management (CRM), customer retention stands out as a pivotal element for sustained growth and competitiveness. The advent of Customer Churn Prediction (CCP) models has been a consequential development in this context, seeking to identify customers at risk of churning in the future (Burez & Van Den Poel, 2007). However, the conventional evaluation metrics applied to these models, such as the Area Under the Receiver Operator Curve (AUC) and top-decile lift, often fall short of aligning with the overarching business goal of maximizing customer retention campaign profit (Verbeke et al., 2012).

This research contends that a paradigm shift is requisite in CCP methodologies to bridge the gap between theoretical model evaluation and practical business impact. To address this misalignment, the Expected Maximum Profit measure for Customer Churn (EMPC) was introduced, offering a novel performance metric that allows practitioners to select models directly based on their business impact (Verbraken et al., 2013). However, despite these efforts, both conventional and novel algorithms continue to assume that the most profitable retention campaigns result from accurately ranking customers based on their churn propensity, inadvertently overlooking the nuanced distribution of Customer Lifetime Value (CLV) among customers (Janssens et al., 2022).

Recognizing the paramount importance of distinguishing between high-value and low-value customers, recent studies by Janssens et al. (2022) and Óskarsdóttir et al. (2018) have introduced novel metrics and algorithms. The EMPB metric, proposed by Janssens et al. (2022), offers an alternative to profit-driven metrics that fail to account for individual differences in customer value (Verbraken et al., 2013). Furthermore, the B2Boost algorithm, introduced by the same authors, focuses on high-value customers through instance-dependent learning, leading to a significant improvement in campaign profitability. Óskarsdóttir et al. (2018) introduces the EMPC vector, comprising individual EMPC values computed with customer-specific CLV values. Analyzing summary statistics of this vector provides a more nuanced insight into the distribution of EMPC values across the heterogeneous customer group.

Despite the promise of these individualized metrics, their integration within profit-driven models demands considerable effort, often resulting in slower operational pace. This paper explores an alternative, less cumbersome path by advocating the incorporation of pre-processing techniques. Inspired by machine learning fairness literature (Kamiran & Calders, 2012), these techniques, including Massaging, Reweighting, and Resampling, aim to rectify the under-representation of high-value customers within the targeted group.

The significance of this approach lies in its potential to streamline the enhancement of churn campaign profitability. Through adept feature cleaning, engineering, and selection, even simple logistic regression models can compete with sophisticated algorithms (Coussement et al., 2017). Similarly, this research finds a comparable effect when pursuing improved representation of targeted high-value customers. This contribution is two-fold: firstly, it addresses the oversight in the current pre-processing literature by proposing techniques to amplify the often-underrepresented high-value customer fraction within the targeted group (Janssens et al., 2022; Óskarsdóttir et al., 2018). Secondly, by framing the challenge as one involving the sensitivity of high-value customers as a group, this research provides a simpler interface for improving churn campaign profit on par with existing profit-driven techniques. Additionally, it delves into the essential prerequisites for the success of individualized profit-driven techniques, exploring situations where conventional methods are most advisable.

2 References

[1] Burez, J., & Van Den Poel, D. (2007). CRM at a pay-TV company: Using analytical models to reduce customer attrition by targeted marketing for subscription services. Expert Systems with Applications, 32(2), 277–288.

[2] Coussement, K., Lessmann, S., & Verstraeten, G. (2017). A comparative analysis of data preparation algorithms for customer churn prediction: A case study in the telecommunication industry. Decision Support Systems, 95, 27–36.

[3] Janssens, B., Bogaert, M., Bagué, A., & Van Den Poel, D. (2022). B2Boost: Instance-dependent profit-driven modelling of B2B churn. Annals of Operations Research.

[4] Kamiran, F., & Calders, T. (2012). Data preprocessing techniques for classification without discrimination. Knowledge and Information Systems, 33(1), 1–33.

[5] Óskarsdóttir, M., Baesens, B., & Vanthienen, J. (2018). Profit-Based Model Selection for Customer Retention Using Individual Customer Lifetime Values. Big Data, 6(1), 53–65.

[6] Verbeke, W., Dejaeger, K., Martens, D., Hur, J., & Baesens, B. (2012). New insights into churn prediction in the telecommunication sector: A profit driven data mining approach. European Journal of Operational Research, 218(1), 211–229.

[7] Verbraken, T., Verbeke, W., & Baesens, B. (2013). A Novel Profit Maximizing Metric for Measuring Classification Performance of Customer Churn Prediction Models. IEEE Transactions on Knowledge and Data Engineering, 25(5), 961–973.

Cheaper neural network ensembles: a case study in manufacturing

A. Thuy
Ghent University, Department of Marketing, Innovation and Organisation
FlandersMake@UGent, core lab CVAMO
arthur.thuy@ugent.be

T. Suys
Delaware BeLux

D. F. Benoit
Ghent University, Department of Marketing, Innovation and Organisation
FlandersMake@UGent, core lab CVAMO

Keywords: Neural network ensembles, Uncertainty quantification, Out-of-distribution data

Introduction

Ensembles of neural networks (NNs), referred to as Deep ensembles, have been regularly used in the operations research (OR) literature. Deep ensembles demonstrate robust performance, both in terms of predictive accuracy and uncertainty quantification (Thuy and Benoit, 2023). Despite their strong performance, applying Deep ensembles in OR poses challenges due to the significant computational and memory requirements. Specifically, these demands increase linearly with the ensemble size. Computation-wise, each member requires a separate forward pass during training and testing. Memory-wise, each member requires a copy of the NN weights. Allocating such substantial resources becomes hard to justify, considering both financial constraints and environmental impact.

Fundamental research has proposed NN ensembles that offer reduced computational overhead while maintaining strong predictive performance and uncertainty quality. However, despite their potential, these methods have not been studied in the context of typical OR applications. This study aims to bridge this gap by conducting a comparative analysis involving a single NN, Deep ensemble, and more cost-effective NN ensembles within an OR case study. The objective is to make reliable uncertainty estimates more accessible in OR applications with NNs, as this is critical for out-of-distribution (OOD) data frequently encountered during deployment.

Methodology

For the advanced NN ensembles, we use a Snapshot ensemble (Huang et al., 2017), Batch ensemble (Wen et al., 2020), and Multi-input multi-output ensemble (Havasi et al., 2020). The NN ensembles are compared on a dataset containing images of manufactured parts. Computational and memory costs are recorded during runtime. The models are trained and tested on images of five object classes, representing the in-distribution set. Additionally, the models are tested on an out-of-distribution dataset of three other object classes, simulating instances during deployment where the data distribution differs from the training data distribution. In the OOD setting, robust uncertainty estimates are crucial to raise a warning and reduce misclassifications by bringing a human expert in the loop for uncertain observations.

Results

This study reveals a clear trade-off between computational cost and uncertainty quality. On the one hand, the single NN is confidently wrong when confronted with OOD data. On the other hand, the Deep ensemble exhibits high uncertainty in predictions on the OOD set. Among the advanced NN ensembles, the Snapshot ensemble demonstrates the strongest performance, nearly matching the Deep ensemble, closely followed by the Batch ensemble. However, the Multi-input multi-output ensemble yields disappointing results, marginally surpassing the single NN. Regarding in-distribution accuracy, the Batch ensemble and Multi-input multi-output ensemble perform comparably to the Deep ensemble, while the Snapshot ensemble performs weaker and is even outperformed by the single NN.

Overall, this work demonstrates that Batch ensembles provide a cost-effective and competitive alternative to Deep ensembles. They perform similarly on the in-distribution set and are slightly less effective on the OOD set while exhibiting a training time speedup of 2X, a test time speedup of 7X, and 5X memory savings. We anticipate that the accessibility of these cheaper NN ensembles will prove advantageous within the OR domain.

References

M. Havasi, R. Jenatton, S. Fort, J. Z. Liu, J. Snoek, B. Lakshminarayanan, A. M. Dai, and D. Tran. Training independent subnetworks for robust prediction. *arXiv preprint arXiv:2010.06610*, 2020.

G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017.

A. Thuy and D. F. Benoit. Explainability through uncertainty: Trustworthy decision-making with neural networks. *European Journal of Operational Research*, 2023.

Y. Wen, D. Tran, and J. Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*, 2020.

Prosumer: a flexible optimization tool to support businesses on their decarbonization journey

M. Boucquey
Engie Impact, Advanced Analytics
matthieu.boucquey@engie.com

Keywords: MILP, Generation Expansion Planning

To mitigate climate change as far as possible, it is urgent to dramatically reduce our greenhouse gas emissions. For this reason, public and private entities make an effort to join their forces to initiate a disruptive and large-scale energy transition. However, even with growing citizen pressure and financial incentives for decarbonization, business actors as well as policymakers are still experiencing major difficulties in meeting their emission commitments. They are even often distraught when trying to fully engage in their decarbonization journey. This is mainly due to the new challenges and the generally increasing complexity of the energy landscape coming from intermittency caused by the injection of renewable energy sources, decentralization of the energy system blurring the roles of producer and consumer, interdependency of energy systems and land use. Nevertheless, business actors seek to better understand possible energy levers (i.a., green electricity, hydrogen, e-fuels, biogas,...) and technology choices to accelerate their decarbonization journey to be able to fulfil their commitments without compromising their business profitability.

To support the Engie group and its clients on this journey, a web-based optimization software, called Prosumer, has been developed by ENGIE Impact R&D lab, which is composed of researchers in mathematics as well as experts in energy systems. As its name

suggests, Prosumer (Producer-consumer) targets to harmonize decarbonization targets with the production potential and the energy needs of a given business actor while providing financially optimal solutions. Finding the most appropriate and profitable option constitutes a complex problem, as it depends on the amount of particularities and requirements to decarbonize, involving a large number of decisions (i.a., investment decisions, production decisions) to make. Therefore, Prosumer utilizes a tailored Mixed Integer Linear Program (MILP) to describe the underlying generation expansion planning problem which is solved using a state-of-the-art commercial solver. A linear Generation Expansion Planning (GEP) problem consists of finding the optimal energy equipment mix to meet future demand. The GEP problem is expressed as commodity and technology agnostic, involving flexible space and time dimensions. While the nodes and edges of the flow network, parameters for the different assets, the accessible commodity markets, and the loads to satisfy are defined by the modelling choices, investment and operational decisions are the outcome of the mathematical optimization, leading to the lowest total cost of ownership. Prosumer, therefore, provide business developers with great flexibility to describe the future portfolio of assets and used commodities.

While also other modelling frameworks have similar mathematical approaches, Prosumer outperforms many approaches thanks to the accessibility of the framework (graphical web-interface is provided), the scalability (cloud-based solving), and the user experience (users do not need to be mathematical modelers), increasing the efficiency of the day-to-day business of Engie Impact consultants. To this end, Prosumer finds application, e.g. to perform Net Zero Factory studies, but is also used by other Engie segments as a pre-feasibility tool to dimension heat or cold networks, green energy supply for datacenters robust to crisis scenarios, charging infrastructure for electric mobility or even equipment to produce green hydrogen with a renewable energy portfolio.

The presentation will involve practical business examples where the full modelling potential Prosumer and the concrete realizations and achievements following its optimal investment recommendations will be showcased.

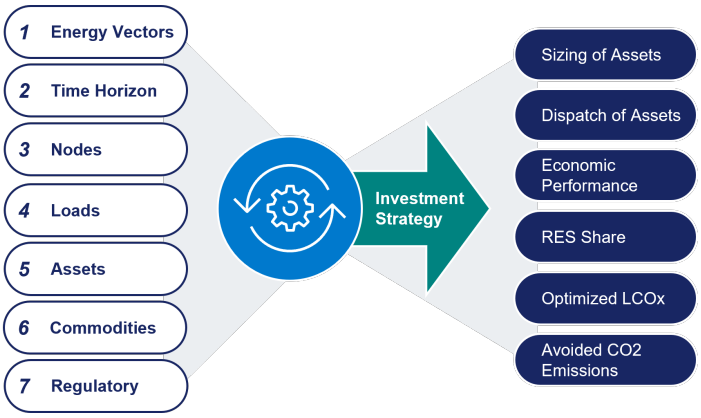

Figure 1: Illustrative diagram of the input modelling potential and the optimized output information of Prosumer.

Navigating Uncertainties for Ensuring Energy Security

L. Mentens
University of Antwerp
lander.mentens@student.uantwerpen.be

J. Braet, A. Herremans and J. Springael
University of Antwerp, Department Engineering Management
Faculty of Business and Economics
Universiteit Antwerpen
Prinsstraat 13, B-2000 Antwerpen, Belgium

Keywords: Energy security, Linear optimization modeling, Economic Dispatch Problem, Crossborder capacity

The expansion of generation capacity in Europe is witnessing substantial growth, primarily driven by the adoption of clean energy sources that are poised to transform global electricity systems. However, the rapid growth in capacity presents challenges due to the variable and uncontrollable nature of renewable energy, leading to a less predictable and more uncertain energy supply. This paper will address the issue of supply security by examining five distinct scenarios for Belgium's future energy mix. This is done by using the urbs software, which is a linear programming model specifically tailored to handle complex energy systems, considering both total system cost and total emitted CO_2 of the modeled system. Each scenario entails different iterations to examine the impact of changing parameters. The analysis of the planned actions reveals risks for blackouts, primarily due to the lower baseload capacity of nuclear energy and the increasing share of renewable energy sources in Belgium and neighboring countries. The results show that adequate cross border capacity and a diversified renewable energy mix supplemented with storage capacity can mitigate these risks for potential blackouts. It proves to be essential for maximizing renewable energy integration and reduces the reliance on conventional, higher-emitting power sources leading to environmental benefits and lower economic costs. This paper highlights the relationship between cross border capacity, storage capacity, CO_2 emissions, and costs. It underscores the importance of proactive planning and investment in cross border capacity and storage infrastructure to enhance the resilience and decarbonization of energy systems. By optimizing the utilization of renewable energy sources and minimizing environmental impacts, the findings contribute to the ongoing transition towards a sustainable and secure energy supply.

Application of Reinforcement Learning for the Optimization of Sustainable Energy Systems

G. Keutgen
Engie Impact, Advanced Analytics
guillaume.keutgen@engie.com

Keywords: Reinforcement Learning, Optimization, Energy

1 Curse of uncertainty in sustainable energy optimization

In the contemporary context, the imperative of effecting an energy transition has emerged as a pivotal element in the pursuit of a sustainable global future. Nations and industrial sectors are progressively incorporating an expanding proportion of sustainable energy sources within their energy portfolios. Nevertheless, the integration and effective management of these emergent technologies to align with energy demand introduce a heightened level of operational complexity. Industrial operators now contend with stochastic and weather-dependent energy productions, an expanded spectrum of technologies and the intricacies associated with storage and energy conversion management.

In response to these challenges, the adoption of digital tools powered by mathematical modeling and optimization methodologies assumes a pivotal role. Operational decisions of sustainable assets are improved by these tools. They allow to formulate robust, secure, and economically viable strategies.

Nevertheless, a significant challenge in addressing optimization problems related to sustainable energy management systems lies in their inherent **stochastic** and **non-linear** characteristics.

- **Stochastic data**: solar farms and wind turbines production, electricity and hydrogen prices
- **Non-linear behavior**: charge and discharge efficiency of batteries, efficiency of energy conversion in hydrogen electrolysers

These aspects pose considerable difficulty when employing conventional optimization approaches. In many cases, there is no hope to meet operational time constraints. Recognizing this challenge, Engie Impact is presently directing a segment of its research efforts towards comparative analyses between optimization techniques and reinforcement learning (RL) methodologies. RL methods are trained to acquire optimal decision-making

insights from simulations and involve the construction of a probabilistic surrogate model for the underlying problem. These methods diverge from the conventional optimization practices avoiding to re-start the computation from scratch each time case parameters are changed. Moreover, the exploration capacities of RL seem of potential to deal with non-linearities.

2 Perspectives on industrial applications

Following an academic literature review, we can see more and more examples of successful applications of RL in energy management systems. This gives the wish to evaluate whether these successes could be replicated on our own industrial applications.

Engie Impact's current research in this domain concentrates on two aspects of energy management: battery and electrolyser dispatch. Both scenarios include a renewable source with uncertain production levels. Moreover, the system is interconnected with the electricity grid and subject to uncertain market prices. The battery dispatch challenge is to efficiently store and release energy with a non-linear charging and discharging efficiency. Similarly, in the second case, electrolyser non-linear production and H_2 storage should be optimized. In both cases, the goal is to maximize revenue from energy sales while ensuring a robust response to cope with stochastic input data.

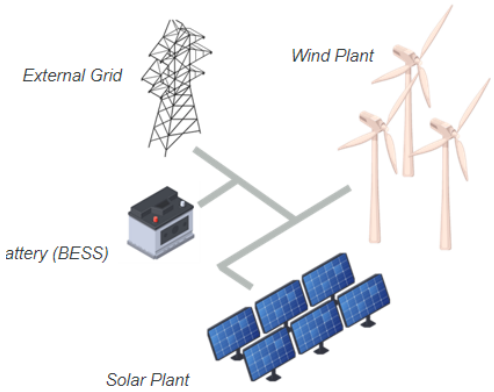

Figure 1: Battery dispatch problem configuration

In this presentation the two cases will be outlined as well as preliminary results and learned insights.

HIST-critical graphs and Malkevitch's conjecture

J. Goedgebeur
KU Leuven Campus Kulak Kortrijk, Department of Computer Science
jan.goedgebeur@kuleuven.be

K. Noguchi
Tokyo University of Science, Department of Information Sciences
noguchi@rs.tus.ac.jp

J. Renders
KU Leuven Campus Kulak Kortrijk, Department of Computer Science
jarne.renders@kuleuven.be

C. T. Zamfirescu
Ghent University, Department of Applied Mathematics, Computer Science and Statistics
czamfirescu@gmail.com

Keywords: HISTs, Spanning Trees, Exhaustive Enumeration, Hamiltonian Cycle

In a given graph, a spanning tree without vertices of degree 2 is called a *HIST* (*homeomorphically irreducible spanning tree*). In the theory of these spanning trees, *HIST-free graphs*, i.e., graphs not containing a HIST, play an important role. Many fundamental questions remain unanswered.

We investigate HIST-freeness from two perspectives. Firstly, we focus on *HIST-critical* graphs, i.e., HIST-free graphs in which every vertex-deleted subgraph contains a HIST. Secondly, we study a conjecture by Malkevitch [2] stating that planar 4-connected graphs must contain a HIST.

This topic is closely related to the study of hamiltonian graphs and hamiltonicity-related concepts. A graph is *hamiltonian* if it has a cycle visiting every vertex of the graph. In other words, a hamiltonian cycle is a connected spanning subgraph in which *every* vertex has degree 2. A HIST is a connected spanning subgraph in which *no* vertex has degree 2. Malkevitch's conjecture is a HIST-analogue of the famous theorem by Tutte [4] stating that planar 4-connected graphs are hamiltonian and HIST-critical graphs are an analogue of the well-studied concept of *hypohamiltonian graphs*, i.e., non-hamiltonian graphs in which every vertex-deleted subgraph is hamiltonian.

To study these concepts, we implemented an efficient branch and bound algorithm for counting the number of HISTs in a given graph. It can then also be used to determine

HIST-freeness or HIST-criticality of a graph. We use this algorithm to obtain computational results and to aid in the proving of theoretical results.

In particular, we performed an exhaustive search for the smallest HIST-critical graphs of at least a given *girth*, i.e., the length of a smallest cycle. See Figure 1 for the three smallest examples. Note that these have girth equal to 3. Some of the graphs found in this way gave rise to an infinite family of HIST-critical graphs. This family together with the exhaustive search are then used to prove the existence of HIST-critical graphs for nearly every order.

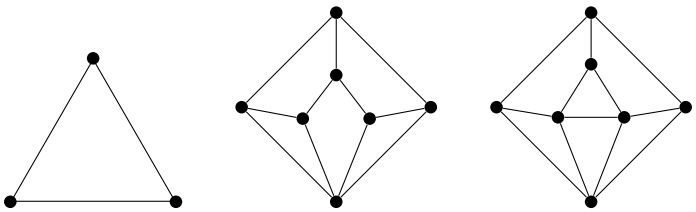

Figure 1: The three smallest HIST-critical graphs.

We performed similar computations for the planar case and also here found an infinite family of planar HIST-critical graphs (see Figure 2), mirroring the study in the hypohamiltonian case, where the non-existence of planar hypohamiltonian graphs was conjectured by Grünbaum [1], but later an infinite family of them was given by Thomassen [3].

Using our algorithm we confirmed Malkevitch's conjecture up to order 21 and look at the 4-connected planar graphs having the fewest number of HISTs for each order. It seems that for even orders this minimum is attained by the antiprism graphs and we prove a formula giving the number of HISTs in such graphs.

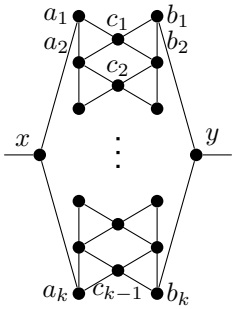

Figure 2: An infinite family of planar HIST-critical graphs.

References

[1] B. Grünbaum. Vertices missed by longest paths or circuits. *J. Comb. Theory Ser. A.* **17** (1974) 31–38.

[2] J. Malkevitch. Spanning trees in polytopal graphs. *Annals of the New York Academy of Sciences*, **319** (1979) 362–367.

[3] C. Thomassen. Planar and infinite hypohamiltonian and hypotraceable graphs. *Discrete Math.* **14** (1976) 377–389.

[4] W.T. Tutte. A theorem on planar graphs. *Trans. Amer. Math. Soc.* **82** (1956) 99–116.

Computational bounds for the minimum order of an edge-girth-regular graph

J. Jooken
Department of Computer Science
KU Leuven Campus Kulak-Kortrijk
jorik.jooken@kuleuven.be

J. Goedgebeur
Department of Computer Science
KU Leuven Campus Kulak-Kortrijk
Department of Applied Mathematics, Computer Science and Statistics
Ghent University
jan.goedgebeur@kuleuven.be

Keywords: Graph algorithms, Extremal problems, Edge-girth-regular graphs

In 2018, Jajcay, Kiss and Miklavič [1] introduced a class of highly symmetrical graphs, called *edge-girth-regular graphs* (abbreviated as *egr* graphs). For integers v, k, g and λ the authors define an $egr(v, k, g, \lambda)$ graph G as a graph consisting of v vertices such that every vertex has degree k, the shortest cycle in G consists of g edges and every edge of G is contained in precisely λ shortest cycles. This class of graphs contains a number of famous graphs that often appear as the minimizers or maximizers of certain functions on graphs. Fig. 1 shows the Petersen graph, which is an $egr(10, 3, 5, 4)$ graph that appears for example in Cambie, Goedgebeur and Jooken [2] as the graph containing the most number of connected induced subgraphs among all 3-regular graphs on at most 10 vertices.

In this work, we are interested in determining $n(k, g, \lambda)$: the minimum number of vertices v such that an $egr(v, k, g, \lambda)$ graph exists for given parameters k, g and λ (or ∞ if no such graph exists). To this end, we develop an exhaustive graph generation algorithm for enumerating all $egr(v, k, g, \lambda)$ graphs based on efficient pruning rules that we integrate in a backtracking algorithm. We implemented this algorithm and conducted a computational experiment on a cluster, which took about 6 CPU-years. This experiment resulted in new extremal graphs and improvements to several lower and upper bounds from the literature for $n(k, g, \lambda)$.

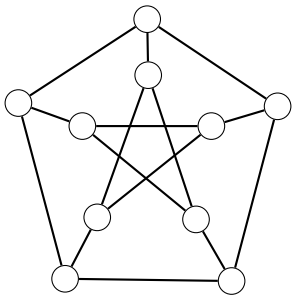

Figure 1: The Petersen graph.

References

[1] Jajcay, Robert, Kiss, György, and Miklavič, Štefko. Edge-girth-regular graphs. *European Journal of Combinatorics* 72 (2018): 70–82.

[2] Cambie, Stijn, Goedgebeur, Jan, and Jooken, Jorik. The maximum number of connected sets in regular graphs. arXiv preprint arXiv:2311.00075 (2023).

Comparative Analysis of Graph Randomization: Methods, Pitfalls, and Best Practices

B. De Clerck
Royal Military Academy,
Department of Mathematics
Ghent University,
Department of Economics
bart.declerck@ugent.be

F. Van Utterbeeck
Royal Military Academy,
Department of Mathematics

L. E.C. Rocha
Ghent University,
Department of Economics
Department of Physics and Astronomy

Keywords: Graph randomisation, Null models, Network science

1 Context

Graph randomization is pivotal in network science for distinguishing significant structural patterns from those arising by chance. While random graphs have been a subject of study since the inception of graph theory, the choice of randomization method for specific research questions is not always clear, especially for non-mathematicians. Moreover, practitioners often rely on existing tools, which may not transparently convey their assumptions and algorithms. As graph randomization becomes increasingly popular for assessing the statistical significance of network metrics and topologies, understanding these methods' assumptions and implications is crucial.

This work explores the intricacies of graph randomization, its potential pitfalls, and the various methods and tools available. We aim to provide a comprehensive overview of different randomization approaches, their algorithms, and the tools implementing them. Additionally, we highlight the importance of method selection for particular research questions and raise awareness of the risks associated with improper randomization, offering guidance to avoid common pitfalls.

2 Methods

We compare conceptual approaches to graph randomization, such as microcanonical and canonical ensembles, and examine the algorithms they employ, including edge swaps,

rewiring, and entropy maximization. We also discuss different metrics for quantifying statistical significance in network structures. Our analysis uses a range of graphs and metrics to demonstrate the coherence, or lack thereof, across various tools and methods.

3 Results

Our findings indicate significant discrepancies in the conclusions drawn from different tools and methods. We provide insights into the often opaque assumptions of these methods and offer recommendations for selecting the most suitable approach for specific research questions. These insights are valuable for network science researchers and practitioners, offering a critical perspective on the methodologies and tools used to analyze complex real-world systems.

4 References

[1] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. "Random graphs with arbitrary degree distributions and their applications". In: Phys. Rev. E 64 (2 July 2001), p. 026118

[2] Tiziano Squartini. Maximum-entropy networks : pattern detection, network reconstruction and graph combinatorics. Cham, Switzerland: Springer, 2017. isbn: 978-3-319-69436-8.

[3] Nicolò Vallarano et al. "Fast and scalable likelihood maximization for Exponential Random Graph Models with local constraints". In: Scientific Reports 11.1 (2021).

[4] Fabien Viger and Matthieu Latapy. "Efficient and simple generation of random simple connected graphs with prescribed degree sequence". In: J. Complex Networks 4.1 (2016), pp. 15-37.

Leveraging Decision Trees to Optimize Medicine Stockpile Policies

S. Lemmens, O. Fourmentraux, H. de Vries
Rotterdam School of Management, Technology and Operations Management Department
s.lemmens@rsm.nl

T. Breugem
Tilburg University, School of Economics and Management

Keywords: Medicine shortages, stockpiling, decision trees

1 Introduction

Drug shortages are recognized as a growing problem worldwide. It has been declared a major public health concern by the European Parliament and Council, and it has been qualified a "national security issue" by U.S. experts and has prompted the creation of numerous task forces. As such, it is studied under multiple lights to identify the causes, consequences, and possible mitigation strategies.

Drug shortages lead to increased patient monitoring, delayed care, increased length of hospitalization, readmission due to adverse events, operations and care cancellation, or in some rare cases, death. Economic consequences are also substantial because of the time pharmacy teams spend dealing with shortages and the cost of purchasing more expensive alternatives.

We know that a myriad of factors can contribute to a drug shortage, such as the availability of raw materials, manufacturing problems, regulatory issues, the structure of the tendering process, just-in-time inventory, demand variability, or various other disruptions along the pharmaceutical supply chain.

Medicine stockpiling has been recommended multiple times, and it is being discussed or implemented in many countries. Stockpiling involves keeping a minimum stock level specified in agreements, guidelines, and/or legislation. Such policies vary from country to country regarding the number of months of demand to keep in stock, the subset of medicines it applies to, and who is responsible for managing and storing it.

2 Research problem

While strategic stockpiles have been argued to have the potential to reduce shortages significantly, this measure is heavily debated. First, it does not address the underlying issues and vulnerability of the supply chain. Second, stockpiles are very expensive and can lead to much waste because of the medicine's limited shelf life. Third, countries typically express the quantity to be stockpiled as a certain number of months of demand, and this number is the same across all medicines included in the stockpile. Hence, there is little to no tailoring of stockpile quantities, while medicines differ greatly in terms of price, demand fluctuation, holding cost, supply chain risks, or criticality for health. Additionally, a significant challenge in implementation is the difficulty in reaching consensus on the location for inventory storage, as ample storage space is required. Calculating the optimal inventory levels according to various factors could prevent hoarding and overstocking, hence reducing the overall cost.

Tailoring is needed, but it increases the complexity of the policy. An effective public policy must remain simple to promote understanding, ease of implementation, and good governance. Furthermore, complex policies may require more resources and be difficult to enforce, while simple policies are easier to understand and require fewer resources to implement. Decision trees are an elegant way to deal with the trade-off between tailoring and policy complexity, and governments often use them for this reason. Policymakers face a decision problem where they must determine the optimal way to split the tree, how to perform it, and which policy or course of action should be assigned to each leaf node. While models and corresponding exact methods for decision tree optimization exist, they disregard the latter aspect of this problem. Instead, the models are mostly designed for classification or regression purposes.

3 Methodology

This paper proposes a new mixed-integer optimization model that addresses the trade-off between tailoring and complexity in policy optimization. Building on the mixed-integer optimization formulation for classification decision trees brought by Bertsimas and Dunn (2017), we revisit the objective function to go beyond misclassification minimization and demonstrate that doing so allows us to optimize policies. Adding a new set of constraints to the optimal classification trees model was also necessary to allow splits on categorical values which frequently appear in policy optimization.

We use our model to optimize the inventory levels of medicines stockpiled in The Netherlands and devise a cost-effective policy that is easy to implement. The model generates a set of splits that aim to minimize the cost of the stockpile inventory in terms of holding and stock-out costs. We estimate the stock-out costs from data on shortages and demand for medicines over the past decade in The Netherlands. The decision tree is designed to split over data on registered medicines and their medicinal characteristics, such as their Anatomical Therapeutic Chemical (ATC) code, route of administration, dosage, form, etc. The inventory level of each medicine can be determined by classifying them according to the characteristics and levels indicated by the tree.

A Novel Multi-Objective Model for Data-Driven Scattered Storage Assignment in Warehouses!

M. Bahadornia
Research Group Logistics, Hasselt University, Belgium
mostafa.bahadornia@uhasselt.be

K. Ramaekers
Research Group Logistics, Hasselt University, Belgium
katrien.ramaekers@uhasselt.be

K. Braekers
Research Group Logistics, Hasselt University, Belgium
kris.braekers@uhasselt.be

T. Cornelissens
ANT/OR Research Group, University of Antwerp, Belgium
trijntje.cornelissens@uantwerpen.be

Keywords: Warehouse Management, Data-Driven, Multi-Objective Model

Today's competitive retail market is faced with high expectations from customers on both delivery (time) and price (cost). This situation has led retailers such as Amazon and Zalando to improve their replenishment and picking strategies at their distribution centres to achieve a higher degree of efficiency [1]. One of these replenishment strategies is Scattered Storage Assignment (SSA). The underlying idea behind SSA is to unbundle each received Stock-Keeping Unit (SKU) and spread it over different storage positions in the warehouse. As the distance travelled by pickers to retrieve items is a crucial issue in warehouse operations management, SSA increases the average adjacency of pickers to SKUs, irrespective of his/her actual position, which leads to less travel distance in the warehouse [2].

To the extent of our review, the prevalent definition and measure used in literature for scatteredness is based on [3]. We present a new data-driven scatteredness measure which extends the concept of scatteredness to include more real-world requirements and takes customer order data into account. To do so, a novel multi-objective mathematical model is proposed in this study. This model assigns SKUs to storage positions with the aim to (a) maximize the scatteredness; (b) minimize splitting order-lines by aiming to collect all items of an order-line from the same location; (c) maximize order correlation of items

close to each other by reducing the distance between items which are frequently ordered together and (d) maximize adjacency of frequently-ordered items to the depot.

The proposed model in this study has several benefits. Firstly, as order frequency of items is not identical, their degree of scatteredness is weighted differently. Therefore, association rules which describe customer order behaviour, are included in the proposed scatteredness measure. Secondly, as mentioned by [4], balanced dispersion of each SKU through the warehouse is important. In other words, if 18 units of an SKU are located in three different locations with the inventory of (6, 6, 6), it is more balanced than (5, 1, 12). Consequently, this characteristic is also embedded in the proposed measure. Thirdly, pairwise distance between various locations of an SKU is taken into account. Namely, if 2 units of an SKU are located in two different locations where their pairwise distance is 30 meters, they are more scattered than the situation where their pairwise distance is 5 meters. Fourthly, in contrast to existing SSA measures that allow a single SKU in each position, the proposed measure in this study allows multiple SKUs in a position which can have different spatial capacities. Fifthly, previous research requires a pre-defined degree of scatteredness [5], while our proposed multi-objective SSA model determines the optimal storage location without needing pre-calculation of a fitted scatteredness degree. Like this, the provided approach in this study not only scatters the inventory fit to the context of the corresponding business but also tries to keep the inventory level of each SKU in each location in a way that minimizes splitting order-lines.

In order to avoid having to find the Pareto frontier, which is complex and may require long computation times, lower and upper bounds for each of the four objectives are calculated. These bounds are then applied to re-scale each objective, such that they can be used in a weighted-sum approach. Next, a data-driven method for assigning the weight of each objective is introduced in this study. Finally, as the proposed model is both non-linear and non-differentiable, a meta-heuristic solution algorithm based on Differential Evolution is developed.

References

[1] Boysen, N., R. De Koster, and F. Weidinger, Warehousing in the e-commerce era: A survey. European Journal of Operational Research, 2019. 277(2): p. 396-411.

[2] Weidinger, F. A precious mess: on the scattered storage assignment problem. in Operations Research Proceedings 2016: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), Helmut Schmidt University Hamburg, Germany, August 30-September 2, 2016. 2018. Springer.

[3] Weidinger, F., Picker routing in rectangular mixed shelves warehouses. Computers and Operations Research, 2018. 95: p. 139-150.

[4] Pawar, N.S., S.S. Rao, and G.K. Adil, A New Measure for Scattering of Stocks in E-commerce Warehouses. IFAC-PapersOnLine, 2022. 55(10): p. 1357-1362.

[5] Albán, H.M.G., T. Cornelissens, and K. Sörensen, Scattered storage assignment: Mathematical model and valid inequalities to optimize the intra-order item distances. Computers and Operations Research, 2023. 149: p. 106022.

Stacking inbound sheets with scattered arrivals

B. Raa, S. Van Volsem
Ghent University, Dpt. Industrial Systems Engineering and Product Design
birger.raa@ugent.be, sofie.vanvolsem@ugent.be

Keywords: Stack loading, Integer programming, Heuristics

1 Problem description

We study an inbound stacking problem at a company that assembles (large) constructions from steel sheets. The steel sheets are produced by a supplier and delivered to the assembly company in batches. An incoming delivery batch usually contains a mix of sheets for different construction projects. For efficient floor space utilization in the warehouse, the steel sheets are stacked. Once assembly for a certain project is started, all corresponding sheets are collected from the warehouse. In fact, most projects require multiple days of assembly work and are split up into sub-projects that correspond to one working day in assembly. Only one project is active on any given day, and the planning of which project is assembled on which day is known beforehand.

The stacks of incoming steel sheets need to be built up in such a way that sheets required for an earlier project are not put below sheets for a later project. Otherwise, the sheet needs to be 'excavated' from the stack, resulting in additional workload for the cranes collecting the sheets. Furthermore, there are stacking constraints related to height, length and width. The stack height is limited to a maximum (e.g., 2m). For stability reasons, a larger sheet can only be put on top of a smaller one if the relative difference in length and width is less than a certain fraction t (t is for tolerance). E.g. if $t = 10\%$, then a sheet that is put on top of a 6.00m by 1.20m sheet cannot be longer than 6.60m or wider than 1.32m.) Finally, because of the limited carrying capacity of the warehouse floor, there is also a weight constraint for each stack: the weight is limited to the surface of the sheet at the bottom of the stack, multiplied with the floor carrying capacity (e.g., 10 ton per m^2).

The supplier produces the sheets in batches based on the required metallurgic characteristics and dimensions (length, width, thickness). As a result, the delivery of sheets required for a specific project is scattered over time. Some sheets already arrive four weeks before the project assembly starts, while others may only arrive a few days beforehand. This means the stacking problem is crucial, since the number of available slots for building stacks in the warehouse is limited.

2 Solution approaches

To gain some initial insights, we have developed, implemented and evaluated two algorithms. The first algorithm is a construction heuristic that builds a solution by adding the sheets one by one to an existing or additional stack. It is a greedy heuristic that assigns sheets to stacks without taking into account the other sheets that still need to be added. The major advantage of this heuristic is that the stack logic embedded in it is easy to understand, which is often an important consideration for practitioners. The disadvantage is that this heuristic will often produce solutions with too many stacks.

The second algorithm is based on a MIP formulation of the problem that is solved using Gurobi. To expedite finding the optimal solution, the solution from the heuristic is used as a warm start, and symmetry-breaking constraints are added on top of the functional constraints. This algorithm consists of two phases. In the first phase, the number of stacks is minimized. However, there are often multiple solutions with the same number of stacks. Therefore, a second phase is added, in which the number of stacks is constrained to the minimum found in the first phase. The modified objective function in phase two is then used to minimize the cumulative 'lifetime' of the generated stacks. This leads to a better grouping of sheets in stacks according to their arrival and retrieval times, so that stacks are created later and/or emptied earlier. This ensures that physical stack spaces are occupied for a shorter duration and are thus freed up more quickly to store new stacks. The significant advantage of this MIP-based algorithm is that it provides a solution with a guaranteed minimum number of stacks (under the given constraints). The drawback is that the computational time can quickly become prohibitive because the search space increases combinatorially with the number of sheets in the dataset. Another disadvantage is that the resulting solutions are not built based on an easily explainable stacking logic and may therefore sometimes seem counterintuitive.

3 Results

Computational results based on data provided by the company provide insights into the trade-off between computational efficiency and solution quality and show the added value of the 2-phase exact approach.

In further research, we plan to develop a hybrid MIP-based heuristic solution approach that efficiently trades off runtime with solution quality. Also, we will study how transforming some of the stacking constraints into soft constraints (e.g., allowing sheet excavations at a penalty cost) affects both the solution methodology and the solution quality. Furthermore, we will validate our solution approach by testing it on problem variants found in the literature.

What characteristics define a Good Solution in Social Influence Minimization Problems?

I. Lozano-Osorio, J. Sánchez-Oro and A. Duarte
Universidad Rey Juan Carlos, Madrid, Spain
isaac.lozano@urjc.es,jesus.sanchezoro@urjc.es,abraham.duarte@urjc.es

K. Sörensen
University of Antwerp, Antwerp, Belgium.
kenneth.sorensen@uantwerpen.be

Keywords: Problem knowledge, Social Network Influence and Heuristics.

1 Introduction

Social Network Influence Problems aim to strategically select users to optimize a certain criterion, such as maximizing or minimizing influence, that can be restricted for a certain budget. Some real-work applications are: rumors, fake news, misinformation, etc.

This paper addresses a $\mathcal{NP}$-hard variant of the Influence Minimization Problem (IMP) [2] which, given a Social Network (SN) and a set of Malicious Nodes (MN), aims to minimize the influence propagation by selecting a set of Block Users (nodes that block the propagation). To our knowledge, the best algorithm proposed to solve this problem is named Greedy Replace (GR), which is based on dominator trees [2].

Formally, a SN is represented by a graph $G = (V, E)$, where the set of vertices V corresponds to users, and the set of edges E indicate the relationships between users in the SN. The objective of IMP is to strategically place a set of Blockers (B) that are devoted to block the propagation of misinformation ($B \subseteq V \setminus MN$), where $|B| = b$), to minimize the number of activated or influenced nodes.

2 From data to problem-specific knowledge

The selection of key features that characterize a good solution is a critical part of this research [1]. The objective of this research is to understand the features that define a good solution.

The identification of such distinguishing properties could help in finding those good solutions more effectively, which is the goal of all heuristics. The different properties of

a solution (features) are obtained by evaluating different metrics over it, thus requiring to transform the structure of a solution into quantitative features. The generated features will then serve as input to the predictive model. This step is highly exploratory, since there are no guidelines about which metrics should be included, and it is usually problem dependent.

A decision tree was chosen as the classifier, which is able to quickly categorize the data set and determine the importance of each feature. This classifier achieved a prediction accuracy of 99.87%. Performing an individual analysis of the features selected, the average sum of the propagation blockers neighbors and the average sum blockers in-degree differ are the key features for IMP. The proposed heuristic uses these features as a greedy criterion.

3 Computational Results

The testbed of instances used is the same set considered in the previous work, which is derived from the well-known Stanford Network Analysis Project (SNAP).

		GR				g_{in}			
	b	**Avg.**	**Time(s)**	**Dev.**	**#B**	**Avg.**	**Time(s)**	**Dev.**	**#B**
WCM (1)	20	10077.13	28.40	0.02%	7	**10077.02**	**28.34**	**0.00%**	**8**
	40	9750.55	**55.49**	**0.01%**	**7**	**9750.50**	55.80	0.03%	**7**
	60	9513.50	**81.54**	0.34%	6	**9512.33**	81.99	**0.05%**	**7**
	80	9301.72	**106.63**	0.44%	6	**9299.53**	107.71	**0.03%**	**7**
	100	9137.11	**131.14**	0.91%	5	**9132.98**	132.87	**0.00%**	**8**
		9556.00	**80.64**	0.34%	31	**9554.47**	81.34	**0.02%**	**37**
TV (2)	20	**15716.42**	211.56	**0.00%**	**8**	**15716.42**	**199.17**	**0.00%**	**8**
	40	**14979.66**	413.64	**0.00%**	**8**	**14979.66**	**387.90**	**0.00%**	**8**
	60	14465.62	604.81	**0.00%**	**8**	14465.62	**569.66**	**0.00%**	**8**
	80	13411.32	781.76	0.01%	7	**13411.30**	**739.35**	**0.00%**	**8**
	100	**13004.12**	952.28	**0.00%**	**8**	13004.21	**899.81**	0.01%	7
		14315.43	592.81	**0.00%**	**39**	14315.44	**559.18**	**0.00%**	**39**

Table 1: A comparison between state-of-the-art method and a heuristic approach based on the best features identified by a supervised algorithm.

Table 1 shows competitive results when comparing both approaches with two diffusion models (WCM and TV). In terms of computing time, both are similar and by analyzing the number of best solutions found, g_{in} is able to reach 76 while GR obtains 70 of the 80 instances.

Acknowledgments

The authors acknowledge the support of the Spanish Ministry of Ciencia, Innovación MCIN/AEI/ 10.13039/501100011033/FEDER, UE) under grant ref. PID2021-126605NB-I00 and PID2021-125709OA-C22.

References

[1] Arnold, F., Sörensen, K.: What makes a VRP solution good? the generation of problem-specific knowledge for heuristics. Computers & Operations Research **106**, 280–288 (Jun 2019)

[2] Xie, J., Zhang, F., Wang, K., Lin, X., Zhang, W.: Minimizing the influence of misinformation via vertex blocking (2023)

Modeling Sand Ripple in Mine Countermeasure Simulations with Stochastic Optimal Control

P. Blondeel
Royal Military Academy, Belgium
Department of mathematics
philippe.blondeel@mil.be

F. Van Utterbeeck
Royal Military Academy, Belgium
Department of mathematics
filip.vanutterbeeck@mil.be

B. Lauwens
Royal Military Academy, Belgium
Department of mathematics
ben.lauwens@mil.be

Keywords: Stochastic Optimal Control, Mine Countermeasures, Sand Ripples

1 Introduction

Modeling and simulating mine countermeasures (MCM) search missions performed by autonomous vehicles is a challenging endeavor. In this work, we present a novel way to model dunes present on the bottom of the ocean which are called sand ripples. These ripples greatly impact the sea mine detection capabilities of the autonomous vehicles, see [4]. In order to account for these ripples in our simulations, we model them in a stochastic optimal control framework. After which, we use the framework to calculate the trajectories of autonomous vehicles in a zone where sea mines are suspected to be present.

2 Methodology

We continue from the work of [1]. In our work, we implemented the optimization problem in the Julia **InfiniteOpt.jl** package, see [2]. Contrary to [1], we formulated the optimization problem such that the total mission time is minimized for a given probability of non-detection, i.e.,

$$\min T_f, \tag{1}$$

subjected to

$$\min \mathbb{E}[q\left(T_F\right)] := \min \int_{\Omega} \mathrm{e}^{-\int_0^{T_F} \gamma(\boldsymbol{x}(\tau),\omega) d\,\tau} \phi\left(\omega\right) d\,\omega \tag{2}$$

where the position of the autonomous vehicle is given by

$$\boldsymbol{x}\left(t\right) := f(x(t), y(t), \psi(t), r(t)). \tag{3}$$

The position as defined in Eq. (3) is governed by the differential equations given in [1].

Our approach to model the sand ripples consists of multiplying the gamma function, $\gamma\left(\boldsymbol{x}\left(\tau\right),\omega\right)$, of Eq. (2), with a set of functions containing the soft rectangular function as defined in [3]. This approach will force the optimization software to select trajectories, for the autonomous vehicle, that are perpendicular to the position of the sand ripples. Indeed, in order to be able to detect mines in the presence of sand ripples, the autonomous vehicle needs to traverse the sand ripples perpendicularly.

3 Results

The to be surveyed zone consists of a square domain $\Omega = [5, 25]^2$ where sand ripples are present in the left upper triangle at a 135 degree angle, see Fig. 1. The trajectory for one vehicle is given in Fig. 2 for a requested probability of non-detection of 30 %, a starting point at $(14.5, 15.0)$, and a total mission time of 28.17 sec. Observe that in Fig. 2, at point $(12.5, 15.0)$, the vehicle is traveling at an angle of approximately 135 degrees, meaning that the detection of mines at this location is impossible, i.e., the color in front of the vehicle stays blue.

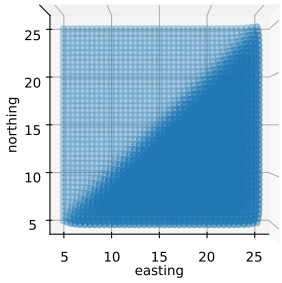

Figure 1: The upper left triangle has sand ripples at a 135 degree angle, while the lower right triangle has no sand ripples.

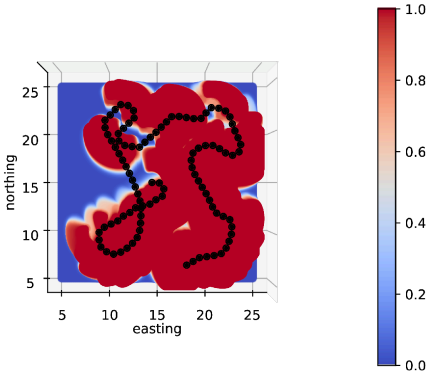

Figure 2: Trajectory for one vehicle for a requested probability of non detection of 30%. The red area denotes the zone that has been surveyed.

4 References

[1] Sean Kragelund, Claire Walton, Isaac Kaminer, and Vladimir Dobrokhodov. Generalized optimal control for autonomous mine countermeasures missions. *IEEE Journal of Oceanic Engineering*, 46(2):466–496, 2021.

[2] Joshua L. Pulsipher, Weiqi Zhang, Tyler J. Hongisto, and Victor M. Zavala. A unifying modeling abstraction for infinite-dimensional optimization. *Computers & Chemical Engineering*, 156, 2022.

[3] Ido Schaefer, Hillel Tal-Ezer, and Ronnie Kosloff. Semi-global approach for propagation of the time-dependent schrödinger equation for time-dependent and nonlinear problems. *Journal of Computational Physics*, 343:368–413, 2017.

[4] Veronika Yordanova, Bart Gips, Thomas Furfaro, and Samantha Dugelay. Coverage path planning for mine countermeasures: Adapting track orientation. In *OCEANS 2019 - Marseille*, pages 1–7, 2019.

Determining stability intervals for intra-critreria parameters in PROMETHEE II

A. Flachs
Université libre de Bruxelles (Belgium), CoDE-SMG
alexandre.flachs@ulb.be

Y. De Smet
Université libre de Bruxelles (Belgium), CoDE-SMG
yves.de.smet@ulb.be

Keywords: Multi-criteria decision aid, Promethee II ranking, sensitivity analysis

In multiple criteria decision aid, one usually distinguishes three main approaches [5]: aggregating, interactive and outranking methods. PROMETHEE [2] belongs to the latter family. These methods have been used in many practical applications [1] and have been subject to multiple methodological extensions. Among these methods, PROMETHEE II allows building a complete ranking of the alternatives. Its structure can be decomposed into three parts:

1. A normalisation of the difference between each pair of alternatives for each criterion;
2. A weighted sum of these normalized differences over all criteria;
3. The definition of a preference structure based on net flow scores.

This is performed by taking into considerations both intra and inter criteria preference parameters. Of course, the study of the sensitivity of the ranking regarding the input parameters is central for the assessment of the method's outputs reliability. To that end, Mareschal and Doan & De Smet studied the stability of PROMETHEE II rankings regarding the inter-criteria parameters (referred to as weights) [4, 3]. The authors proposed exact methods to compute *intervals of stability* on weights which preserve a ranking. Their approach is based on linear programming.

To the best of our knowledge, there has not been any study on the stability of PROMETHEE II ranking regarding the intra-criteria parameters. In this case, a major difficulty resides in the fact that the relation between the results of PROMETHEE II and the intra-criteria parameters is in general not linear and does not respect easily manageable mathematical properties. In this work, we propose an exact sequential algorithm to compute *regions of stability* in which the relative ranking is preserved under some restrictive hypotheses

regarding the used normalisation. The results of the algorithm are presented and studied on different problem instances.

The proposed method is expected to be extendable to relax these hypotheses. For example to preserve a complete ranking, or to preserve relative ranking when changing multiple parameters simultaneously. We propose different research directions to investigate these questions.

References

[1] Majid Behzadian, R.B. Kazemzadeh, A. Albadvi, and M. Aghdasi. Promethee: A comprehensive literature review on methodologies and applications. *European journal of operational research*, 200(1):198–215, 2010.

[2] Jean-Pierre Brans and Yves De Smet. *PROMETHEE Methods*, pages 187–219. Springer New York, New York, NY, 2016.

[3] N.A.V. Doan and Y. De Smet. An alternative weight sensitivity analysis for promethee ii rankings. *Omega*, 80:166–174, 2018.

[4] Bertrand Mareschal. Weight stability intervals in multicriteria decision aid. *European Journal of Operational Research*, 33(1):54–64, 1988.

[5] Philippe Vincke. *Multicriteria decision-aid*. J. Wiley, Chichester, 1992.

Evaluating the PROMETHEE II ranking quality

B. Coquelet
Université libre de Bruxelles (Belgium), CoDE-SMG
boris.coquelet@ulb.be

G. Dejaegere
Université libre de Bruxelles (Belgium), CoDE-DesLab
Gilles.dejaegere@ulb.be

Y. De Smet
Université libre de Bruxelles (Belgium), CoDE-SMG
Yves.de.smet@ulb.be

Keywords: PROMETHEE II, Quality index, Ranking Consistency

Multi-Criteria Decision Aid (MCDA) aims to help solve problems consisting of alternatives that are simultaneously evaluated on conflicting criteria. Usually, three main problem settings are considered; choosing, sorting or ranking alternatives. In MCDA, three prominent families of methods can be identified: aggregating, outranking and interactive procedures. This study will focus on PROMETHEE which belongs to the outranking methods (Brans & Vincke [1]). It is based on pairwise comparisons to generate a partial or a complete ranking of alternatives. As attested by Behzadian et al. [2] and Mareschal [3], they are widely used in practice and have been subject to extensions to sorting, clustering, group decision-making, etc. Within this family, PROMETHEE I produces a partial ranking (leaving some pairs of alternatives incomparable to each other) and PROMETHEE II produces a complete ranking.

In MCDA, the method, as well as the parameters, are chosen by a decision-maker. The selection of exact values for these parameters, while often complex for a decision-maker, is of the outermost importance. However, it is inherently difficult to assess whether illogical rankings originate from a poor choice of parameters, inadequate ranking method or that the data itself is not suitable for ranking. In the Analytical Hierarchical Process method (AHP), this aspect is evaluated through what is called an inconsistency index (Saaty [4]). However, to the best of the authors' knowledge, such considerations have never been made in the application of the PROMETHEE II method.

Thus, this contribution aims to address this shortcoming by providing an assessment of the resulting ranking quality given by PROMETHEE II. This is achieved by evaluating the conversion from what is called the preference matrix to the uni-dimensional ranking. Indeed, a quality index can be computed from this preference matrix which can indicate the amount of inconsistencies that are then neglected when computing the final complete ranking.

References

[1] J.-P. Brans, *L'ingénierie de la décision: l'élaboration d'instruments d'aide a la décision*. Université Laval, Faculté des sciences de l'administration, 1982.

[2] M. Behzadian, R. Kazemzadeh, A. Albadvi, and M. Aghdasi, "Promethee: A comprehensive literature review on methodologies and applications," *European Journal of Operational Research*, vol. 200, no. 1, pp. 198–215, 2010.

[3] B. Mareschal. "The Promethee bibliographical database." (2020), [Online]. Available: `http://en.promethee-gaia.net/bibliographical-database.html` (visited on 08/26/2022).

[4] T. L. Saaty, "The analytic hierarchy and analytic network processes for the measurement of intangible criteria and for decision-making," in *Multiple Criteria Decision Analysis: State of the Art Surveys*, S. Greco, M. Ehrgott, and J. R. Figueira, Eds. Springer New York, 2016, pp. 363–419.

Disjunctive scheduling using interval decision variables with Hexaly Optimizer

L. Blaise
Hexaly
lblaise@hexaly.com

Keywords: disjunctive scheduling, solver, modeling

1 Introduction

The aim of this talk is to show how the rich set-based modeling formalism of Hexaly Optimizer (formerly LocalSolver) can be used to efficiently model many industrial scheduling problems, using generic operators only. More specifically, we will focus on the advantages brought by the combined use of interval and list variables for modeling disjunctive scheduling problems. We will also show how the solver exploits this modeling to provide quality solutions quickly, through local search moves and a solution repair algorithm based on constraint propagation.

Hexaly Optimizer is a mathematical optimization solver based on various operational research techniques, combining exact methods such as linear, nonlinear and constraint programming, and heuristic methods such as local search [2]. Its algorithms deliver excellent performance on scheduling problems: 2.2% deviation from the best known solution on the Job Shop Problem with up to 2000 tasks, 0.3% on the Flexible Job Shop Problem with up to 500 tasks, 1.4% on the Resource-Constrained Project Scheduling Problem with up to 300 tasks, all within one minute of computation.

2 List and interval-based modeling

The presence of disjunctive resources in a scheduling problem is characterized by non-overlap constraints on tasks assigned to the same resource. A simple way of writing this constraint consists in exploiting the order of the tasks: each task can only start after the previous task has finished.

With Hexaly Optimizer's modeling formalism, this constraint is written using two types of decision variables. On the one hand, each resource is represented by a list variable[1] : the elements of the list correspond to the tasks executed on this resource, arranged in

[1] In Hexaly Optimizer's modeling formalism, a list variable of domain size n is a decision variable whose value is a permutation of a subset of $\{0, ..., n-1\}$.

ascending order of execution dates. In addition, interval decision variables are used to represent the execution range of each task.

```
order <- list(nbTasks);
task[0...nbTasks] <- interval(0, horizon);
```

We then write the non-overlap constraint using a variadic "and" operator. This constraint reads "for any position i in the list, the task at position i must be executed before the task at position $i+1$". This formulation based on list and interval variables has the advantage of allowing the user to write the non-overlap constraint in $O(n)$ only, where n represents the number of tasks, as opposed to $O(n^2)$ when considering tasks two by two and using only interval variables.

```
constraint and(0...nbTasks-1, i => task[order[i]] < task[order[i+1]]);      // (1)
```

3 Small-neighborhood search

In addition to providing a straightforward way to model non-overlap constraints, expression (1) gives the model a strong structure that can be exploited by the solver, particularly in local search movements.

Hexaly Optimizer's local search component thus implements various scheduling-specific movements, which are automatically activated or deactivated according to the characteristics of resources and tasks (fixed or variable duration, for instance). Examples include: inserting a new task on a machine, merging two neighboring tasks, exchanging resources between two tasks, exchanging the start dates of k tasks assigned to the same resource, splitting a task...

These movements exploit the interaction between variables representing task execution ranges and their order on the machines. For example, for the resource change move, we remove the index of the chosen task from the list of its initial machine and add it to the list of another machine at the correct position, thus ensuring that the tasks remain sorted by ascending start date.

4 Solution repair by constraint propagation

However, small-neighborhood search alone does not enable the solver to achieve good performance on disjunctive scheduling problems. Indeed, in a good solution to such a problem, the precedence and non-overlap constraints between tasks are very tight. Moving from a good solution to a slightly better one then requires changing the execution ranges of a large number of tasks.

To overcome this problem, we implemented a solution repair algorithm based on constraint propagation [1]. The principle of this algorithm is as follows. When the solution obtained after a local move is infeasible, it is gradually repaired, one constraint at a time. The goal is to reach feasibility by extending a promising but infeasible move, rather than canceling it. The constraints are therefore repaired by modifying variable values always in the same direction. For example, if a task has already been moved back, either during

the move or during a previous repair, it can be moved back further, but cannot be moved forward.

5 References

[1] L. Blaise, C. Artigues, and T. Benoist. Solution Repair by Inequality Network Propagation in LocalSolver. In *Parallel Problem Solving from Nature – PPSN XVI*, 2020.

[2] F. Gardi, T. Benoist, J. Darlay, B. Estellon, and R. Megel. *Mathematical Programming Solver Based on Local Search.* John Wiley & Sons, Ltd, 2014.

A Constraint Programming Approach for Aircraft Disassembly Scheduling

C. Thomas
UCLouvain, ICTEAM
charles.thomas@uclouvain.be

P. Schaus
UCLouvain, ICTEAM
pierre.schaus@uclouvain.be

Keywords: Aircraft Dismantling, Scheduling, Constraint Programming

1 Introduction

This research is done as part of the Planum project which consists in studying and developing technologies and tools in order to facilitate the recycling of end-of-life aircrafts. The problem studied in this paper concerns the scheduling of the operations taking part in the dismantling phase of the recycling process, from the reception of the plane to the sectional cutting and shredding of the carcass. Additionally, technicians have to be assigned to tasks during the scheduling and some tasks require specific certification levels. The main objective is to minimize the total time taken by the whole extraction process. This is modeled with a *makespan* value that corresponds to the time step at which the last operation finishes. A secondary objective is to minimize the dismantling cost by limiting the use of more costly resources. This problem is a variant of the Multi-Skill Project Scheduling Problem (MCPSP) [1], which is itself a variation of the well known Resource Constrained Scheduling Problem (RCPSP) [2].

Additional constraints arise from the specific context of aircraft dismantling: First, some parts of the plane may have space restrictions that limit the number of technicians working at the same time there. Thus, the plane is divided into locations that each have an occupancy limit corresponding to the maximum number of technicians allowed to work there at the same time. Second, the plane must be kept balanced during the whole disassembly process by ensuring that the difference of mass between its extremities does not overstep given thresholds.

2 Model

The approach used to tackle this problem uses Constraint Programming (CP). The model is implemented in CP Optimizer [3] and relies on conditional time-intervals [4] which are used to represent the tasks to perform as well as the assignation of technicians. Task variables are always set to present while assignation variables are optional. Some technicians may be unavailable at specific times during the scheduling period. These unavailabilities

are also modelled as interval variables. All the optional assignation and unavailability interval variables of a same technician are added to a sequence variable which is subject to a *noOverlap* constraint that ensures that a technician is never assigned to more than one operation simultaneously and is not assigned when unavailable.

Balance and occupancy constraints are modelled using cumulative functions. There are two cumulative functions used for the balance constraints: one that represents the difference of mass between the front and rear of the aircraft and one does the same for left and right wings. When weight is removed in a balance zone as part of a task, it is either added to or subtracted from the relevant cumulative function. Occupancy constraints also use cumulative functions: For each location in the airplane, a cumulative function models the number of technicians working in this location. This cumulative function is linked to the operation activities taking place at this location and must not overstep the capacity of the location.

An *Alternative* constraint is used to link the task activities with the corresponding optional assignation activities. Finally, precedence constraints ensure that preceding activities are finished when an activity starts. The two objectives of the problem are solved using a lexicographical search: First, the makespan objective is solved to optimality or until a given limit is reached. Second, the cost objective is minimized subject to an additional constraint that prevents the makespan objective to regress.

The model is assessed using a set of scenarios comprising up to 1500 tasks, which are derived from real data provided by an industrial partner. Our experiments demonstrate that the model can effectively identify feasible solutions for all instances. However, proving optimality is only feasible for instances with a smaller scale. For larger instances, solutions are found with a gap of 15% to the lower bound. The model, instances and results are available at the following repository: `https://github.com/cftmthomas/AircraftDisassemblyScheduling`.

Acknowledgments This work was funded by the Walloon Region (Belgium) as part of the Planum project. We thank Sabena Engineering for allowing the diffusion of an anonymized version of the dataset provided.

References

[1] Odile Bellenguez and Emmanuel Néron. Lower bounds for the multi-skill project scheduling problem with hierarchical levels of skills. In *International conference on the practice and theory of automated timetabling*, pages 229–243. Springer, 2004.

[2] Sönke Hartmann and Dirk Briskorn. An updated survey of variants and extensions of the resource-constrained project scheduling problem. *European Journal of Operational Research*, 297(1):1–14, 2022.

[3] Philippe Laborie, Jérôme Rogerie, Paul Shaw, and Petr Vilím. Ibm ilog cp optimizer for scheduling: 20+ years of scheduling with constraints at ibm/ilog. *Constraints*, 23:210–250, 2018.

[4] Philippe Laborie, Jérôme Rogerie, Paul Shaw, Petr Vilím, and Ferenc Katai. Interval-based language for modeling scheduling problems: An extension to constraint programming. *Algebraic Modeling Systems: Modeling and Solving Real World Optimization Problems*, pages 111–143, 2012.

Lock scheduling with non-identical parallel chambers

M. Davari
Skema Business School, France,
morteza.davari@skema.ed

M. Ranjbar
Ferdowsi University of Mashhad, Iran

D. Briskorn
University of Wuppertal, Germany

Keywords: Lock scheduling, polynomial-time algorithm, time-indexed formulation.

1 Problem definition

Rivers and waterways, in addition to road, rail, and air transport, serve as natural infrastructures well-suited for freight transport. Despite being an age-old method, transporting goods over inland waterways proves reliable, cost-effective, and environmentally friendly within logistics and supply chains. According to the European Commission, the energy consumption for water transport is approximately 17% of that for road transport and 50% of rail transport. Eurostat reported in 2014 that waterways contribute to around 12.3% of Germany's total freight transport infrastructure, and this figure is approximately 6.6% for the European Union (EU).

A lockage or lock movement denotes a singular operation of a chamber, encompassing the entry of one or more ships into the chamber, the adjustment of water levels from downstream to upstream or vice versa, and the subsequent exit of the ships from the chamber. Note that ships approach from either the upstream or downstream side. The term *lockage time* corresponds to the duration required to execute this operation. The chamber's capacity signifies the maximum number of ships that can be concurrently accommodated during a single lockage. Locks are comprised of either a single chamber or multiple parallel chambers. For example, the Wijnegem lock, situated in Belgium, linking the Albert canal to the Antwerp port, comprises three non-identical parallel chambers.

Formally, we consider a single lock consisting of non-identical parallel chambers. The chambers operate independently of each other, and each is characterized by two numbers: a length of the chamber and a lockage time (or simply a duration). Moreover, a set of non-identical ships arrive on either side of the lock. Each ship is characterized by an arrival time, a length, a waiting cost per unit time and a position, either downstream or upstream. We aim to assign ships to chambers and schedule each lockage of chambers

such that the total waiting cost is minimized.

2 Complexity

Following a straightforward reduction from 3-partition, we conclude the problem is strongly NP-hard. Also, we show the complexity of the problem when ships are identical for an arbitrary number of chambers is strongly NP-hard since its no-wait counterpart is already known to be strongly NP-complete. Nevertheless, we provide a polynomial-time algorithm to solve this special case when the number of chambers is constant.

Theorem 1 *Assuming a constant number of chambers, the special case with identical ships is solvable in* $\mathrm{O}(n^{m+3})$ *time.*

3 Initial results

We propose different formulations to solve the problem. For the time-indexed formulation, we propose additional inequalities enforcing the following property.

Property 1 *Each chamber movement either directly follows upon a previous movement or starts upon some* a_i *while containing ship* i*.*

Preliminary findings indicate that the time-indexed formulation, incorporating additional inequalities, exhibits effectiveness in solving instances of problems involving up to 100 jobs. The absence of these additional inequalities, however, significantly hampers the performance of the time-indexed formulation. For the special case with identical ships and for small m, an early implementation of the proposed polynomial algorithm for Theroem 1 demonstrates a slightly superior performance compared to the time-indexed formulation. For the case where $m = 1$, a dedicated iterative dynamic programming approach is performant. Additionally, early results suggest that an associated set-covering formulation may not yield computational advantages, while Bender's decomposition approach exhibits some potential.

4 References

[1] Passchyn, W., Briskorn, D., and Spieksma, F., 2019, "No-wait scheduling for locks", *INFORMS Journal on Computing*, Vol. 31(3), pp. 413-428.

[2] Passchyn, W., Coene, S., Briskorn, D., Hurink, J. L., Spieksma, F. C., and Berghe, G. V., 2016, "The lock-master's problem", *European Journal of Operational Research*, Vol. 251(2), pp. 432-441.

Project scheduling with an external resource

L. Gallois
Skema Business School, Lille and KU Leuven, Faculty of Economics and Business
leonie.gallois@kuleuven.be

M. Davari
Skema Business School, Lille
morteza.davari@skema.edu

J. Matuschke
KU Leuven, Faculty of Economics and Business
jannik.matuschke@kuleuven.be

Keywords: project scheduling, Benders decomposition, outsourcing

1 Introduction

Outsourcing is a popular means to achieve flexibility and reduce the cost of ownership in many industries (Briskorn et al, 2021). Both in project and machine scheduling, outsourcing can take the form of external rented resources. Determining the length of the rental periods for such resources is a challenging problem.

In this paper, we aim to schedule a set of normal jobs and a set of (external) resource jobs within a precedence-constrained project as defined by Brucker et al (1999). Each job has a processing time and induces a cost depending on its completion time. The external resource also induces a renting cost per unit time. The external resource will be present at the site from the start of the first resource job until the end of the last resource job. We consider a discrete-time horizon representing the scheduling horizon.

The jobs have precedence constraints, meaning that each job can only start when all its predecessors are finished. At any point in time, we can process an infinite number of regular jobs but only one resource job.

We consider a very generic objective function in which we minimize the sum of the costs, consisting of a scheduling cost and a renting cost. The scheduling cost can be any piecewise linear function such as the *weighted total completion time* and the *makespan*. Even for these well-known objectives, the problem is strongly NP-hard as shown in Briskorn et al (2021).

2 Benders algorithm

A time-indexed formulation of the problem includes binary decision variables that define the order based on which the resource jobs are scheduled and binary time-indexed

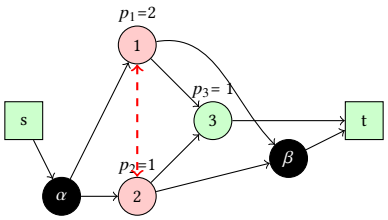

Figure 1: Disjunctive arcs for resource jobs

decision variables that decide the final schedule. In such a formulation, only one block of constraints requires both types of variables, while the other block requires only one variable type. This allows us to devise a Benders decomposition approach to solve the problem. Since the jobs needing the resource cannot be processed simultaneously, the master problem fixes the order of the resource jobs by setting variables that define the direction of the associated disjunctive arcs, see Figure 1. Once these variables are fixed, the sub-problem schedules the jobs by solving a problem whose dual is equivalent to a min-cost flow problem.

3 Results

Proposition 1. *When the solution to the master problem is integral, the dual of the sub-problem reduces to a max flow problem.*

Conjecture 1. *Considering a non-decreasing in the completion time of the job cost function, the dual of the sub-problem reduces to a shortest path problem which is solved in polynomial time with the Dijkstra algorithm.*

This is the first time we can reduce a project scheduling problem to a shortest path problem according to our literature review.

References

[1] Briskorn, D., Davari, M., Matuschke, J., 2021. Single-machine scheduling with an external resource. European Journal of Operational Research 293, 457–468.

[2] Artigues, C., 2017. On the strength of time-indexed formulations for the resource-constrained project scheduling problem. Operations Research Letters 45, 154–159.

[3] Möhring, R., Schulz, A., Stork, F., Uetz, M., 2001. On project scheduling with irregular starting time costs. Operations Research Letters 28, 149–154.

[4] R.H. Mohring, A.S. Schulz, F. Stork and M. Uetz (2003). Solving project scheduling problems by minimum cut computations. Management Science. 49(3):330-350.

[5] Brucker, P., Drexl, A., Mo, R., Pesch, E., 1999. Resource-constrained project scheduling: Notation, classification, models, and methods. European Journal of Operational Research.

[6] Papadakos, N., 2008. Practical enhancements to the Magnanti–Wong method. Operations Research Letters 36, 444–449.

Systematic Derivation of a Global Constraint for Routing Optimization

R. De Landtsheer
Combinatorial Algorithmics, CETIC research center, Charleroi
renaud.delandtsheer@cetic.be

Keywords: Local Search, Global Constraints, Formal Methods

1 Global Constraints Based on Transfer Functions

Global constraints are a very efficient class of algorithms that can be used in local search, notably for routing optimization problems. Global constraints provide speed improvements compared to naive implementations because they perform some form of symbolic reasoning on the structure of the problem and exploit mathematical properties such as associativity and commutativity. They might however be time-consuming to develop.

The OscaR.cbls framework has a generic global constraint based on so-called transfer functions, which are used to label segments of routes [2, 3]. A *transfer function* F is a function $T \rightarrow T$ that can be either queried, as any function, or composed with another transfer function $F \circ F$. The transfer function and the composition function together form a monoïd: the composition of two transfer function must be a transfer function. Both F and T are constraint-specific.

The generic global constraint precomputes all transfer functions involving the first and last node of a route. It also precomputes the transfer function in a logarithmic progression to label segments linking positions (0, 2),(2,4),(4,6),(6,8), ...; (0,4),(4,8),(8,16), ...(0,8),(8,16), ...etc. If the composition operator takes $O(1)$-time, this precomputation takes $O(n)$-time on a route of length n. On a 3-opt move, which creates four segments of the initial route, the transfer function of the first and last segments are available, and the transfer function of the two segments in the middle of the route are assembled out of $O(\log n)$ precomputed transfer functions. If the transfer function are evaluated in constant time, the constraint is updated in $O(\log n)$-time.

To implement a new constraint based on this generic global constraint, one must define a transfer function and its associated composition operator. Although the framework hides away a lot of technicalities, this remaining task is challenging.

2 Systematic Derivation of Transfer Function and Composition Operator

To find the transfer function and the composition operator, one can follow an iterative ”trial and error” approach, supported either by proof or by testing. Alternatively, one can also follow a constructive approach, where the definition of the transfer function and composition operator are derived from the definition of the constraint. Derivation approaches are commonly found in formal methods such as B [1]. We could deploy such an approach on a complex constraint.

The considered constraint is a weighted tardiness constraint with early line. This constraint declares two classes of nodes: the first class takes part in a weighted tardiness; the other class has an early line. Each node the of the first class has an associated weight. All nodes the of the second class share the same early line. For this constraint, we could derive the composition operator from the formal definition of the constraint.

We first crafted a transfer function for a single node. It is parameterized by the class of the node, its tardiness weight, and the early line. The transfer function has two inputs: the arrival time at the node and the summed tardiness of the preceding nodes; and two outputs: the leave time from the node and the summed tardiness of this node and its preceding nodes. Then, a suitable transfer function for segments of route was crafted. It enriches the definition of the per-node transfer function with additional parameters. This enriched formula is -unfortunately- an educated guess.

The main task was to define the value of these parameters. This amounts to defining the composition operators. This is where the derivation process took place. We developed the formula obtained by chaining two transfer function algebraically. The resulting large formula was then re-arranged through distributiveness and associativity into the defined pattern of transfer function. Finally, we identified the terms of the resulting formula to the parameters of the transfer function. This defined the composition operator and proved that the transfer function and the composition operator together form a monoïd.

The derivation was performed in a single day. No subsequent demonstration was needed since the formula is correct by construction, and all tests passed immediately. This advocates that the derivation approach was relevant and cost efficient in this case.

Acknowledgement This work is supported by the Service Public de Wallonie Recherche under grant n°2010235 - TRAIL/ARIAC.

References

[1] J.-R. Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, USA, 1996.

[2] R. De Landtsheer, F. Germeau, T. Fayolle, G. Ospina, and C. Ponsard. *Generic Support for Precomputation-Based Global Routing Constraints in Local Search Optimization*, Heuristics for Optimization and Learning, pages 151–165, Sprigner, 2021.

[3] De Landtsheer, R., and Ponsard, C. (2013) *OscaR.cbls: an open source framework for constraint-based local search.* ORBEL 2013.

The Dial-a-Ride Problem with Limited Pickups per Trip

B. Zhao
KU Leuven, Research Center for Operations Research & Statistics
boshuai.zhao@kuleuven.be

K. Wang
Tsinghua University, School of Vehicle and Mobility
cwangkai@tsinghua.edu.cn

W. Wei
Beijing Jiaotong University, School of Economics and Management
weiwenchao@bjtu.edu.cn

R. Leus
KU Leuven, Research Center for Operations Research & Statistics
roel.leus@kuleuven.be

Keywords: dial-a-ride problem, limited pickups per trip, fragment

The Dial-a-Ride Problem (DARP) is an optimization problem that involves determining optimal routes and schedules for several vehicles to pick up and deliver items at minimum cost. Motivated by real-world carpooling and crowdshipping scenarios, we introduce an additional constraint imposing a maximum number L on the number of pickups per trip. This results in two new variants of the DARP, namely the Multi-vehicle Dial-a-Ride Problem with Limited Pickups per Trip (DARP-LPT) and the Dial-a-Ride Problem with Limited Pickups per Trip (DARP-LPT). In MDARP-LPT, vehicles have unique origins and destinations, while in DARP-LPT, all vehicles share a single fixed depot.

We apply a unified fragment-based method to address both problems, where a fragment is a partial path. Specifically, we extend two formulations from Rist & Forbes [1]: the Fragment Flow Formulation (FFF) and the Fragment Assignment Formulation (FAF). Our polyhedral analysis establishes FFF's superiority over FAF, a conclusion further validated by computational experiments. Furthermore, our results show that FFF and FAF significantly outperform traditional arc-based formulations in terms of solution quality and time.

Alyasiry et al. [2] propose full fragments (FF) as longer partial paths, while Rist & Forbes [1] introduce restricted fragments (RF) as shorter partial paths. We propose two novel fragment types: path-enumerated full fragments (PFF) and mixed fragments (MF). PFF in-

volves enumerating all paths and decomposing them into fragments, while MF decomposes a part of FF into shorter fragments while preserving the rest. Subsequently, we assess the computational performance of different fragment types using the formulation FFF. The results show that, in MDARP-LPT with a lower value of L, PFF outperforms FF, and FF performs better than RF. Additionally, for DARP-LPT, MF surpasses both RF and FF in computational performance, especially when instances have more flexible time windows and a higher value of L.

References

[1] Yannik Rist and Michael Forbes (2021). A New Formulation for the Dial-a-Ride Problem. Transportation Science, 55, 1113-1135. doi:10.1287/trsc.2021.1044.

[2] Ali Mehsin Alyasiry, Michael Forbes, Michael Bulmer (2019). An exact algorithm for the pickup and delivery problem with time windows and last-in-first-out loading. Transportation Science, 53, 1695-1705. doi:10.1287/trsc.2019.0905.

Delivering efficiency: The parcel locker puzzle

L. Melis
Maastricht University, School of Business and Economics
lissa.melis@uantwerpen.be

A. Ranjbari
Penn State University, Civil and environmental engineering department
ranjbari@psu.edu

V. V. Gayah
Penn State University, Civil and environmental engineering department
vvg104@psu.edu

Keywords: Parcel lockers, Location-routing problem, Metaheuristics

Our study explores the implementation of parcel lockers as a solution for urban last-mile delivery challenges. To effectively implement parcel lockers, decisions must be made regarding the number of lockers, optimal locations, and internal configurations. Existing research has not fully integrated these decisions with vehicle routes and flexible delivery options, i.e., home or nearby parcel locker delivery. To address this gap, our study defines the parcel locker location-configuration-routing problem with flexible deliveries as an optimization challenge. Our solution approach combines a population-based incremental learning algorithm with an embedded adaptive large neighborhood search heuristic, considering both carrier and customer perspectives to achieve a balanced solution. Our objective function takes into account the failed home delivery costs, vehicle operation and usage costs, locker opening costs, as well as the inconvenience cost for customers when a parcel locker delivery is assigned instead of a home delivery. Using an artificial case study based on real-world demand and costs, we investigate three research questions. Firstly, we analyze the impact of customer inconvenience on parcel locker usage. Secondly, we evaluate the influence of flexible parcel locker delivery, compared to strict home delivery or strict closest-to-home parcel locker delivery. Lastly, we examine the effects of adopting multiple possible locker sizes instead of one fixed size. Our analysis finds major carrier's cost savings by adopting flexible parcel locker delivery and the possibility of multiple parcel locker sizes. We also suggest persuading customers to choose parcel locker delivery by offering discounts.

Index

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Z