

**KU LEUVEN**

**ORBEL**

THE BELGIAN OPERATIONAL  
RESEARCH SOCIETY

# ORBEL 40

The 40th conference of the Belgian Operational Research Society

Leuven, February 5–6, 2026

**Booklet of Abstracts**



Copyright © 2026 KU Leuven and ORBEL

PUBLISHED BY KU LEUVEN AND ORBEL, THE BELGIAN OPERATIONS RESEARCH  
SOCIETY

[HTTPS://ORBEL.BE/](https://orbels.be/)

*February 2026*

# Contents

1	Welcome to Leuven .....	5
2	ORBEL .....	7
3	Committees .....	9
4	Practical Information .....	11
5	Sponsors .....	15
6	Schedule at a glance .....	17
7	Plenary talks .....	19
8	Detailed program .....	21

<b>I</b>	<b>Abstracts</b>	
	TA1	39
	TA2	45
	TA3	53
	TA4	61
	TA5	67
	TB1	72
	TB2	80
	TB3	88
	TB4	94
	TB5	102

TC1	108
TC2	112
TC3	118
TC4	124
TC5	130
TD1	136
TD2	142
TD3	148
TD4	149
TD5	154
FA1	158
FA2	165
FA3	171
FA4	178
FA5	186
FB1	194
FB2	201
FB3	209
FB4	217
FB5	225
FC1	233
FC2	240
FC3	247
FC4	248
FC5	256
List of Authors .....	263

## **1. Welcome to Leuven**

Welcome to ORBEL 40, the 40th Annual Conference of the Belgian Operational Research Society (ORBEL)! ORBEL 40 brings together researchers, practitioners, managers, and students working in operational research, analytics, simulation, data science, and related quantitative methods for decision making. This year's meeting offers a vibrant forum to exchange ideas, present new research, and connect with the national and international OR community.

The conference is organized at the main campus of KU Leuven on February 5-6, 2026, in the historic and dynamic university city of Leuven, which is known for its rich academic traditions and lively cultural atmosphere, providing an inspiring backdrop for scientific exchange and collaboration. All sessions take place in the buildings of the Faculty of Economics and Business and Arenberginstituut - STUK.

ORBEL conferences have a long tradition of fostering excellence in quantitative research and decision-making techniques, with invited plenary talks, parallel contributed sessions, and activities that promote interaction across disciplines. ORBEL also serves as a platform for recognising emerging talent and achievements in the field.

We hope your time in Leuven will be both intellectually stimulating and personally enjoyable. Welcome to ORBEL 40!

## 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.



### **3. Committees**

### Organizing Committee

- Leus, Roel
- Vansteenwegen, Pieter
- Verbeke, Wouter

### Scientific Committee

- |  |  |
|--|--|
| <ul style="list-style-type: none"> <li>• Aghezzaf, El-Houssaine</li> <li>• Arda, Yasemin</li> <li>• Beliën, Jeroen</li> <li>• Benoit, Dries</li> <li>• Blondeel, Wouter</li> <li>• Braekers, Kris</li> <li>• Caris, An</li> <li>• Chevalier, Philippe</li> <li>• Coussement, Kristof</li> <li>• Crama, Yves</li> <li>• De Baets, Bernard</li> <li>• De Causmaecker, Patrick</li> <li>• Defryn, Christof</li> <li>• De Smet, Yves</li> <li>• Fortz, Bernard</li> <li>• Goossens, Dries</li> <li>• Janssens, Gerrit</li> <li>• Kunsch, Pierre</li> <li>• Labbé, Martine</li> <li>• Leus, Roel</li> </ul> | <ul style="list-style-type: none"> <li>• Louveaux, Quentin</li> <li>• Melis, Lissa</li> <li>• Mélot, Hadrien</li> <li>• Papadimitriou, Dimitri</li> <li>• Paquay, Célia</li> <li>• Pironet, Thierry</li> <li>• Sartenaer, Annick</li> <li>• Schauss, Pierre</li> <li>• Schyns, Michaël</li> <li>• Spieksma, Frits</li> <li>• Sörensen, Kenneth</li> <li>• Tancrez, Jean-Sébastien</li> <li>• Vandaele, Arnaud</li> <li>• Vanden Berghe, Greet</li> <li>• Vansteenwegen, Pieter</li> <li>• Van Utterbeeck, Filip</li> <li>• Verbeke, Wouter</li> <li>• Wauters, Tony</li> <li>• Wittevrongel, Sabine</li> </ul> |
|--|--|

### Helping Hands Committee

We thank the following people, listed alphabetically, for their help in organizing the conference. Our apologies to anyone we may have overlooked.

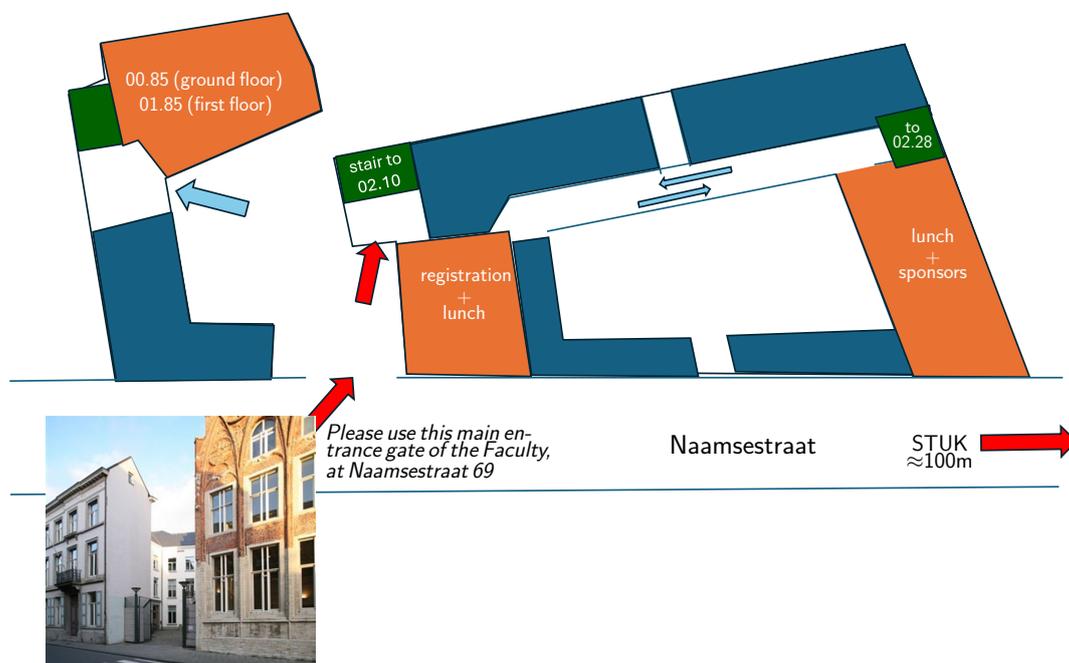
- |  |  |
|--|--|
| <ul style="list-style-type: none"> <li>• Aktaş Dejaegere, Dilay</li> <li>• Bijmens, Ann-Pascale</li> <li>• Caljon, Daan</li> <li>• Danhieux, Leen</li> <li>• De Reymaeker, Kristel</li> <li>• Deprez, Bruno</li> <li>• Geysegom, Ruth</li> <li>• Goethals, Andreas</li> <li>• Jans, An</li> <li>• Kalai, Tarik</li> <li>• Lin, Boyue</li> <li>• Moura, Phablo</li> </ul> | <ul style="list-style-type: none"> <li>• Perneel, Emmeline</li> <li>• Rauh, Felix</li> <li>• Roos, Hilde</li> <li>• Scheepmans, Nico</li> <li>• Seeholzer, Lukas</li> <li>• Van Asten, Saar</li> <li>• Van Belle, Jente</li> <li>• Van Espen, Ann</li> <li>• Van Hoeck, Inneke</li> <li>• Vanderghinste, Emiel</li> <li>• Wens, Maarten</li> <li>• Wu, Yuhang</li> </ul> |
|--|--|

## **4. Practical Information**

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

Rooms are **not** equipped with a computer. Please bring your own laptop with the required connection tool to connect to the HDMI of the projector. If you do not have a laptop, please bring a USB stick to transfer your presentation to the laptop of another speaker in your session. Please arrange this before the session starts. To ensure a timely progression of the session, we ask all speakers to check the functioning of their presentation before the start of the session.

**ORBEL 40** will take place at the [Faculty of Economics and Business](#), located at **Naamsestraat 69**, and we also use a lecture room in STUK (Naamsestraat 96), at 100 meters from the faculty, at the opposite side of the street.



### Conference dinner on February 5 from 19h00

The conference dinner will take place at [De Hoorn – Sluisstraat 79](#), which is at a 20-minute walk from the conference venue. From here, Leuven train station is also approx. 20 minutes on foot.

19:00 – 19:45 Reception (Brouwzaal — standing)

19:45 – 23:00 Conference dinner (Machinezaal — seated, table service)

**Tip for drivers** Enter Parking Vaartkom (Engels Plein 32) in your GPS to reach the restaurant easily and without detours. This way, you will avoid rush-hour stress and will not have to worry about the traffic restriction in front of the building.

**Parking** The city car park Vaartkom is just a 3-minute walk from De Hoorn. The first hour of parking is free, and the maximum daily rate is 6 EUR. A night rate of 3 EUR applies between 7 p.m. and 7 a.m.

**Public transport** At the bus stop *Leuven Sint-Michielskerk*, take bus **2 (Leuven Engels Plein)** and get off at the stop *Leuven OPEK*, or use the [route planner](#) by De Lijn.



## 5. Sponsors

We gratefully acknowledge the generous support of our sponsors.

**Platinum Sponsor**

The logo for hexaly, featuring the word "hexaly" in a lowercase, sans-serif font. The letter "x" is stylized with a light orange or pinkish hue, while the other letters are in a darker orange.

**Gold Sponsors**



## **6. Schedule at a glance**

**Thursday, February 5**

8h30 – ...	<i>Registration</i>	Naamsestraat 69
8h30 – 9h10	Welcome coffee	Kapel
10'	<i>Move to STUK across the street</i>	
9h20 – 9h30	Opening	STUK (other side of the street)
9h30 – 10h30	Plenary Laporte	STUK
10h30 – 11h00	Coffee break	Kapel and registration room
11h00 – 12h20	Parallel sessions TA	STUK, 02.28, 00.85, 01.85, 02.10
12h20 – 13h40	Lunch break	Kapel and registration room
12h20 – 13h40	ORBEL board meeting	00.85
13h40 – 15h00	Parallel sessions TB	STUK, 02.28, 00.85, 01.85, 02.10
15h10 – 16h10	Parallel sessions TC	STUK, 02.28, 00.85, 01.85, 02.10
16h10 – 16h30	Coffee break	Kapel and registration room
16h30 – 17h30	Parallel sessions TD	STUK, 02.28, 00.85, 01.85, 02.10
17h35 – 18h20	ORBEL general assembly	00.85
19h00 – ...	Conference dinner	De Hoorn

**Friday, February 6**

8h45 – ...	<i>Registration</i>	Naamsestraat 69
8h45 – 9h15	Welcome coffee	Kapel and registration room
9h15 – 10h35	Parallel sessions FA	STUK, 02.28, 00.85, 01.85, 02.10
10h35 – 11h00	Coffee break	Kapel and registration room
11h00 – 12h20	Parallel sessions FB	STUK, 02.28, 00.85, 01.85, 02.10
12h20 – 13h20	Lunch break	Kapel and registration room
13h20 – 14h40	Parallel sessions FC	STUK, 02.28, 00.85, 01.85, 02.10
14h40 – 15h10	Coffee break	Kapel and registration room
15h10 – 16h10	Plenary Feuerriegel	STUK
16h10 – 16h30	ORBEL Award announcement and closing	STUK

## **7. Plenary talks**

### **A Short History of the Traveling Salesman and Its Influence on the Field of Combinatorial Optimization**

**Gilbert Laporte**, Professor Emeritus, HEC Montréal, Canada. Professor, School of Management, University of Bath, United Kingdom. Guest Lecturer, Molde University College, Norway.



I will provide a short history of the Traveling Salesman Problem (TSP), starting with the solution of the knight's problem by Leonhard Euler in 1759, and ending with a very efficient implementation of the Lin and Kernighan heuristic by Taillard and Helsgaun in 2019. The talk covers the main exact and heuristic algorithms that have been developed for the undirected and directed cases of the TSP. I will emphasize the influence that the study of the TSP has had on the evolution of the field of combinatorial optimization

### **Causal ML to optimize decisions from observational data**

**Stefan Feuerriegel**, Professor, head of Institute of Artificial Intelligence (AI) in Management, LMU Munich, Germany.



Modern machine learning excels at prediction, yet good predictions do not necessarily lead to good decisions. This talk introduces **causal machine learning** as a framework for learning optimal treatment and action policies directly from data. By estimating the causal effects of alternative decisions, causal ML enables data-driven choices tailored to individual contexts. We illustrate applications across domains: in medicine, selecting personalized treatments that maximize patient outcomes; in business, optimizing pricing, marketing, or intervention strategies; and in the public sector, designing effective policies and resource allocations. Overall, causal ML moves beyond forecasting what will happen to answering the key question for decision-makers: what should we do?

## **8. Detailed program**

**Thursday, 5 February 2026**

- 9h20–9h30*    **Opening session**  
*Room:* STUK 02.C004  
 Welcome by Pieter Vansteenwegen
- 9h30–10h30*    **Plenary talk**  
*Room:* STUK 02.C004  
**Gilbert Laporte**  
 A Short History of the Traveling Salesman and Its Influence on the Field of Combinatorial Optimization  
*Session chair:* Roel Leus
- 11h00–12h20*    **Parallel session TA1**  
*Session chair:* Kristof Coussement  
*Room:* STUK 02.C004
- Computer Vision for Anti-Money Laundering in Time-Varying Transaction Networks  
*Deprez, B., Verdonck, T., Verbeke, W.* ..... 39  
 Robust decision-making for budget allocation in marketing using machine learning  
*Caljon, D., Maarouf, A., Van Belle, J., Verbeke, W., Feuerriegel, S.* ..... 41  
 Entropy-Based Model Evaluation for Guided Contextualized Process Discovery  
*Ahmadi, Z., De Weerd, J., Serral, E.* ..... 43
- 11h00–12h20*    **Parallel session TA2**  
*Session chair:* Frits Spieksma  
*Room:* HOGC 02.28
- Generating Realistic Solar Irradiation Scenarios with Limited Historical Data  
*Hanusa, R., Raa, B., De Vuyst, S.* ..... 45  
 Sensitivity analysis for linear changes of the constraint matrix of a (mixed-integer) linear program  
*Miftari, B., Ernst, D., Louveaux, Q., Derval, G.* ..... 47  
 A Distance Metric for Mixed Integer Programming Instances  
*Maudet, G., Danoy, G.* ..... 49  
 Approximation algorithms for graph search problems with imperfect detection

---

*Van Ee, M., Sitters, R.* ..... 51

11h00–12h20 **Parallel session TA3**

*Session chair:* Hande Yaman

*Room:* HOGM 00.85

DDOLib: A library for Solving Discrete Optimization Problems with Dynamic Programming

*Schaus, P., Meurisse, Q., Kameugne, R., De Landtsheer, R., Germeau, F., Gillard, X.* ..... 53

Sequence Variables: A Constraint Programming Computational Domain for Routing and Sequencing – Extended Abstract

*Delecluse, A., Schaus, P., Van Hentenryck, P.* ..... 55

Integrated location, partitioning and sequencing decisions for efficient power system restoration

*Çalik, H., Van Hertem, D., Yaman, H.* ..... 57

Arc-flow formulation for a two-stage two-dimensional dual bin packing problem for wood reuse

*Bessemans, P., Paquay, C., Dumont, M.* ..... 59

11h00–12h20 **Parallel session TA4**

*Session chair:* An Caris

*Room:* HOGM 01.85

Multi-Commodity Inventory Routing Problem with Pickup and Delivery

*Seeholzer, L., Joubert, J., Braekers, K., Vansteenwegen, P.* ... 61

The optimization of store and warehouse based e-grocery fulfillment

*D’Haen, R., Köhler, C., Campbell, A.M., Ehmke, J.F.* ..... 63

Optimizing Storage Assignment in Temperature Zones for Fruits and Vegetables

*Bahadornia, M., Ramaekers, K., Braekers, K., Cornelissens, T., D’Haen, R.* ..... 64

Real-time railway timetable rescheduling in case of partial blockages

*Van Hoeck, I., Vansteenwegen, P.* ..... 65

11h00–12h20 **Parallel session TA5**

*Session chair:* Dries Benoit

*Room:* HOGS 02.10

Do profit-driven learning methods deliver profit? A policy evaluation study

	<i>Rahman, S., Janssens, B., Rahman, S.</i> .....	67
	A Lightweight Graph Representation for Scalable Deep Learning on 3D CAD Data	
	<i>Van Camp, T., Benoit, D.</i> .....	69
	FoV-Net: Rotation-Invariant 3D Deep Learning for CAD Analysis	
	<i>Ballegeer, M., Benoit, D.</i> .....	71
13h40–15h00	<b>Parallel session TB1</b>	
	<i>Session chair:</i> Matthias Bogaert	
	<i>Room:</i> STUK 02.C004	
	A coevolutionary approach for optimizing interpretable rule ensembles	
	<i>De Lange, S., Bogaert, M., De Bock, K.W., Van den Poel, D.</i>	72
	Trusted Triplet Contrastive Learning for Multi-view clustering	
	<i>Li, Y., Benoit, D.</i> .....	74
	Personalized Advertising with Generative AI through Interpretable User-Creative Clusters	
	<i>Marckx, N., Van Mulders, D., Bogaert, M., Van den Poel, D.</i>	76
	An Empirical Exploration of Explanation-Guided Learning for Causal Effect Estimation	
	<i>Glado, A., Marchese, A., Hirsch, L., Verboven, S.</i> .....	78
13h40–15h00	<b>Parallel session TB2</b>	
	<i>Session chair:</i> David Van Bulck	
	<i>Room:</i> HOGC 02.28	
	POSEIDON: Planned Operations and Scheduling Engine for Integrated Demining in Oceanic Navigation	
	<i>Blondeel, P., Van Utterbeeck, F., Lauwens, B.</i> .....	80
	Scheduling in orbit servicing for mega-constellations	
	<i>Nigudkar, S., Defryn, C.</i> .....	82
	Scheduling a single parallel-batching machine with Run-to-Run constraints	
	<i>Yang, F., Yi, X., Wang, H., Goossens, D., Leus, R.</i> .....	84
	The Unreliable Job Selection and Sequencing Problem	
	<i>Perneel, E., Agnetis, A., Leus, R., Salvadori, I.</i> .....	86
13h40–15h00	<b>Parallel session TB3</b>	
	<i>Session chair:</i> Tony Wauters	
	<i>Room:</i> HOGM 00.85	

A comprehensive analysis of the sparrow-3D heuristic: Bottlenecks in Irregular Strip Packing	
<i>Martens, J., Tollenaere, J.</i> .....	88
Efficient Compliance Checking of EU Driver Regulations within Local Search-Based VRP Solvers	
<i>De Walsche, N.</i> .....	89
An order-first split-second framework for solving the 2D guillotine cutting stock problem with bounded-sized stock	
<i>Hadj Salem, K., Wauters, T.</i> .....	90
Analysis of the relationship between different constraint types in personnel rostering problems	
<i>Garcia Tercero, L.</i> .....	92

*13h40–15h00* **Parallel session TB4**

*Session chair:* Kenneth Sörensen

*Room:* HOGM 01.85

Special cases of the Joint Replenishment Problem (JRP): A mathematical analysis	
<i>Stephan, T., Creemers, S.</i> .....	94
Stochastic bilevel optimization with dynamic decision-making followers	
<i>Jiménez, D., Pagnoncelli, B., Yaman, H.</i> .....	96
Adaptive Hybrid VNS-LNS for Time-Dependent Grocery Delivery Vehicle Routing	
<i>Elsharif, O., Sörensen, K., Galarza Montenegro, B., Caris, A.</i>	98
Contextual optimization with recourse for production-transshipment problem with uncertain yield and demand	
<i>Zong, S., Joubert, J.</i> .....	100

*13h40–15h00* **Parallel session TB5**

*Session chair:* Christof Defryn

*Room:* HOGS 02.10

From Waste to Worth: Extending The Multi-Echelon Network Design Problem Towards Circular Supply Chain Modelling	
<i>Langenaeker, A., Nimmegeers, P., Defryn, C.</i> .....	102
Decision support for afforestation planning: Integrating exact and heuristic spatial optimization methods	
<i>Castillo Reyes, G., Roose, D., Vanden Berghe, G.</i> .....	104
Sustainability Analytics for Assessing Environmental Reporting and Media Perception	

*Diaz Gonzalez Manjarrez, F., Coussement, K., Bouten, L., Meire, M. ....*106

15h10–16h10 **Parallel session TC1**

*Session chair:* Bernard Fortz

*Room:* STUK 02.C004

A Tutorial on Voxel-Based approaches for 3D Irregular Packing Problems

*La Greca, A., Wauters, T. ....* 108

Improved upper bounds for the cage problem via exhaustive generation of lifts

*Stubbe, L., Goedgebeur, J., Jooen, J., Van Den Eede, T., Exoo, G. ....* 109

Exact Methods for the Vehicle Routing Problem with Time Windows and Stochastic Customers

*Han, Y., Yaman, H. ....*111

15h10–16h10 **Parallel session TC2**

*Session chair:* Sabine Wittevrongel

*Room:* HOGC 02.28

A Markovian Mean-Field Model for the Dynamics of Societal Trust  
*Chaudhary, S., De Turck, K., Fiems, D. ....* 112

The Challenge of Quality Assessment in Dynamic Microsimulation  
*Dumont, M., Alsaloum, A., Ernst, J., Weymeirsch, J., Münnich, R. ....* 114

A two-class queueing model for transplant waiting lists with renegeing high-priority patients

*Debevere, T., Devos, A., Wittevrongel, S., Walraevens, J. ...* 116

15h10–16h10 **Parallel session TC3**

*Session chair:* Jean-Sébastien Tancrez

*Room:* HOGM 00.85

First Results on the Impact of Synchronizing Driver Breaks and Charging in Electric Long-Haul Routing

*El Habti, O., Braekers, K., D’Haen, R. ....* 118

Controlling Energy Consumption In Production Systems by Utilizing the Digital Twin

*Khayyati, S., Karabag, O., Berkman, B., Demirtas, S., Tan, B. ....* 120

Assessing the Economic and Environmental Impacts of Peer-to-Peer Crowdshipping in Relation to Spatial Structure <i>Innocente, E., Tancrez, J.</i> .....	122
---	-----

**15h10–16h10 Parallel session TC4**

*Session chair:* Patrick De Causmaecker  
*Room:* HOGM 01.85

New developments in Artelys Knitro 15.1 <i>Dufour, M., Fontan, F.</i> .....	124
Comparing comparative process mining methods: towards a research agenda <i>Riebus, M., Martin, N., Jans, M.</i> .....	126
Federated Routing Optimization for the Physical Internet <i>De Landtsheer, R., Meurisse, Q.</i> .....	128

**15h10–16h10 Parallel session TC5**

*Session chair:* Greet Vanden Berghe  
*Room:* HOGS 02.10

Assessing the Value of Dual Sourcing for Production Continuity <i>De Backker, T., Van Der Haar, J., Boute, R.</i> .....	130
An Integrated Approach to Airport Security Rostering <i>Kamhi, H., Berghman, L.</i> .....	132
Coordinated Scheduling of Equipment and Staff in Blood Component Production <i>Gürsoy, A., Pinto, R., Smet, P., Vanden Berghe, G., Lanzarone, E.</i> .....	134

**16h30–17h30 Parallel session TD1**

*Session chair:* Célia Paquay  
*Room:* STUK 02.C004

Optimizing Tactical Assembly Line Feeding in Picker-to-Part and Part-to-Picker Settings <i>Porbar, G., Verplancke, H., Thanos, E., Limère, V., Aghezzaf, E.</i> 136	
A model-based heuristic for a three-dimensional loading vehicle rout- ing problem with split pickups and time windows <i>Leloup, E., Paquay, C., Pironet, T.</i> .....	138
Component effects of the Late Acceptance Hill Climbing metaheuris- tic	

*Van Den Abeele, M., Sørensen, K.* ..... 140

**16h30–17h30 Parallel session TD2**

*Session chair:* Jente Van Belle

*Room:* HOGC 02.28

Emotion-driven customer satisfaction prediction in call center conversations

*De Cleen, T., Baecke, P., Goedertier, F.* ..... 142

A Bandit-Based Approach to Educational Recommender Systems: Contextual Thompson Sampling for Learner Skill Gain Optimization

*De Kerpel, L., Benoit, D., Thuy, A.* ..... 144

Reinforcement Learning for BESS short-term trading

*Mouffe, M., Brunieau, G., Fontaine, G., Keutgen, G.* ..... 146

**16h30–17h30 Parallel session TD3**

*Session chair:* Quentin Louveaux

*Room:* HOGM 00.85

Invited session: Hexaly tutorial

*Coppé, V., Blaise, L.* ..... 148

**16h30–17h30 Parallel session TD4**

*Session chair:* Pierre Schaus

*Room:* HOGM 01.85

Data-Driven Advanced-Planning-and-Scheduling Implemented in Printed-Circuit-Board Assembly

*Pan, C.* ..... 149

A DP-CP Hybrid Framework for Solving the Job-Shop Scheduling Problem

*Legrand, E., Kameugne, R., Schaus, P.* ..... 150

Collaborative multimodal freight transportation under uncertainty

*Gilet, F., Tancrez, J., Creemers, S.* ..... 152

**16h30–17h30 Parallel session TD5**

*Session chair:* Jeroen Belien

*Room:* HOGS 02.10

Balancing service levels and operational effort in dynamic smart waste collection

---

<i>Aktas Dejaegere, D., Vansteenwegen, P.</i> .....	154
How we scheduled ORBEL: new models, takeaways and feedback	
<i>Rauh, F.</i> .....	156
Generative AI as a Copilot for Mathematical Optimization Modeling	
<i>Lahrs, L.</i> .....	157

**Friday, 6 February 2026***09h15–10h35* **Parallel session FA1***Session chair:* Yasemin Arda*Room:* STUK 02.C004

The Incomplete Traveling Tournament Problem (iTTP)

**Van Bulck, D., Devriesere, K., Goossens, D.** .....158

Novel neighborhood structures for incomplete round robin tournaments

**Devriesere, K., Van Bulck, D., Goossens, D.** ..... 160

Fair Policies for Sequential Resource Allocation

**Wessel, S., Hojny, C., Spieksma, F.** ..... 162

Smart Lotteries in School Choice: Ex-ante Pareto Improvement with Ex-post Stability

**Demeulemeester, T., Aziz, H., Biro, P., Csaji, G.** .....164*09h15–10h35* **Parallel session FA2***Session chair:* Sofie Goethals*Room:* HOGC 02.28

Two-Step Imbalanced PU Learning with Graph Neural Networks

**Bérard, R., Deprez, B., Ortega Vázquez, C., Aerts, M.** ..... 165

ML-Based B2B Quotation Prioritization: Evidence from a Randomized Field Experiment

**Vandendriessche, B., Benoit, D.** ..... 167

On the Difficulty of Detecting Cherry-Picking in Counterfactual Explanations

**Hinns, J., Goethals, S., Martens, D., Van Der Veecken, S., Eugeniou, T.** ..... 168

Consistency in Persona-assigned Large Language Models

**Reusens, M., Baesens, B., Jurgens, D.** ..... 170*09h15–10h35* **Parallel session FA3***Session chair:* Sarah Van der Auweraer*Room:* HOGM 00.85

Leveraging probabilistic imbalance price forecasts for risk-aware trading with battery systems

**Mascarenhas, M., Verstraeten, A., Kazmi, H.** ..... 171

Decision-focused learning for optimal PV-Battery scheduling

**Depoortere, J., Kazmi, H., Driesen, J.** ..... 173

Mitigating the Bullwhip Effect when Forecasting for Inventory Control

**Van Der Haar, J., Dumon, M., Van Belle, J., Boute, R.** ... 174

Exploring Machine and Deep Learning Methods to Forecast Cash Flow in the Logistics Industry

**Wolfe, A., Van Der Auweraer, S., Coussement, K.** ..... 176

*09h15–10h35* **Parallel session FA4**

*Session chair:* Thierry Pironet

*Room:* HOGM 01.85

Enhancing population localization by geo-referencing buildings to census grids

**Kerff, A., Weymeirsch, J., Dumont, M., Vandomme, E.** ..... 178

Improving Pickup and Delivery Routing in a Laundry Service Company

**Moonen, C., Aguayo Quintana, C., Limbourg, S.** ..... 180

Evolution of social acceptability of public lighting reductions throughout the recent energy crisis

**Bebronne, E., Limbourg, S., Cools, M.** ..... 182

Stochastic Revelation in Sequential Decision-Making with a Pharmaceutical Application

**Davari, M., Bodur, M., Pagnoncelli, B.** ..... 184

*09h15–10h35* **Parallel session FA5**

*Session chair:* Emmanouil Thanos

*Room:* HOGS 02.10

Optimizing Urban Logistics by Integrating a Construction Consolidation Center

**Wang, Y., Caris, A., Braekers, K., D'Haen, R.** ..... 186

An application of advanced optimization methods in Non-Emergency Medical Transportation

**Galarza Montenegro, B.** ..... 188

A Large Neighborhood Search Approach for Large-Scale Fleet Assignment Problems

**Annaert, R., Smet, P., Vanden Berghe, G.** ..... 190

A Mixed Integer Programming model for the Strategic Assembly Line Feeding Problem

**Verplancke, H., Thanos, E., Porbar, G., Aghezzaf, E., Limère, V.** 192

11h00–12h20 **Parallel session FB1**

*Session chair:* Dries Goossens

*Room:* STUK 02.C004

- Metaheuristics: beyond the horse race  
*Sörensen, K.* .....194
- Locating a facility in space: A review of cover and facility location problems in space-based applications  
*Deleye, P., Defryn, C.* ..... 195
- What many operations researchers have done wrong and what are the remedies  
*Stützle, T.* ..... 197
- Models and algorithms for the time-dependent traveling salesman problem with drone  
*Zhang, S., Wei, W.* .....199

11h00–12h20 **Parallel session FB2**

*Session chair:* Sam Verboven

*Room:* HOGC 02.28

- A Multi-Agent Approach to Improving XAI-Narratives  
*He, Y., Martens, D.* ..... 201
- Cash or Comfort? How LLMs Value Your Inconvenience  
*Cedro, M., Ichmoukhamedov, T., Goethals, S., He, Y., Hinns, J., Martens, D.* ..... 203
- Foundation Models for Credit Risk Prediction: Game Changer or False Hope?  
*Goethals, A., Baesens, B., Lessmann, S., De Vos, S., Medina-Olivares, V., Martens, D., Mues, C., Verdonck, T., Verbeke, W.*  
 205
- Causal Modeling of Organ Offer Dynamics for Better Allocation  
*Marchese, A., Verboven, S.* ..... 207

11h00–12h20 **Parallel session FB3**

*Session chair:* Hussain Kazmi

*Room:* HOGM 00.85

- Hierarchical forecasting and decision-making: some perspectives from the energy sector  
*Qaiser, T., Mehmood, F., Deconinck, G., Kazmi, H.* ..... 209
- Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand  
*Van Der Auweraer, S., Van Pelt, T., Arts, J.* ..... 211

On the impact of forecast reliability on downstream decision-making <i>Van Belle, J., Verbeke, W., Kourentzes, N.</i> .....	213
Semantic information for variable selection in tactical business forecasting <i>Sagaert, Y., Kourentzes, N.</i> .....	215

11h00–12h20 **Parallel session FB4**

*Session chair:* Jochen De Weerd

*Room:* HOGM 01.85

Unified Multi-Task Predictive Process Monitoring with Transformers and Adaptive Loss Optimization <i>Wuyts, B., vanden Broucke, S., De Weerd, J.</i> .....	217
Causal Estimation of Midday Stem Water Potential for Decision-Oriented Irrigation Scheduling <i>Roa, J., Diels, J., Serral, E.</i> .....	219
SCOPE: Sequential Causal Optimization of Process Interventions <i>De Moor, J., Weytjens, H., De Smedt, J., De Weerd, J.</i> ....	221
Handling Multi-Class Imbalance in Next Activity Prediction for Predictive Process Monitoring <i>He, X., Oyamada, R., De Smedt, J., vanden Broucke, S., De Weerd, J.</i> .....	223

11h00–12h20 **Parallel session FB5**

*Session chair:* Pieter Leyman

*Room:* HOGS 02.10

Parallel Machine Scheduling with Speed Scaling <i>Yang, K., Leus, R.</i> .....	225
A Decomposition Algorithm For the Integrated Flexible Job Shop Scheduling and Operator Timetabling <i>Ghorbani Saber, R., Leyman, P., Aghezzaf, E.</i> .....	226
Efficient Algorithms for Energy-Aware Single-Machine Scheduling with Battery Storage <i>Forte, L., De Vuyst, S., Leyman, P.</i> .....	228
Dynamic Programming for Inspection Station Allocation in Production Lines <i>Hedayati, S., De Vuyst, S., Raa, B.</i> .....	230

13h20–14h40 **Parallel session FC1**

*Session chair:* Stefan Creemers

*Room:* STUK 02.C004

- An integrated approach to operational scheduling in hospital-at-home  
**Maharani, A., François, V., Arda, Y.** ..... 233
- Periodic home healthcare routing with stochastic patients : imitation learning enriched decision-making  
**Tassin, L., François, V., Gendreau, M., Vandomme, E., Arda, Y.** 235
- Merging HIS- and RTLS-data to Strengthen Work Organization Analysis in Nursing  
**Tharwat, H., Martin, N., Depaire, B.** ..... 237
- Integrated planning of public charging infrastructure and low-voltage grid expansion under uncertainty  
**Coulier, B., Çalik, H., Becker, T., Vanden Berghe, G.** ..... 239

*13h20–14h40* **Parallel session FC2**

*Session chair:* Kris Braekers

*Room:* HOGC 02.28

- The Dial-a-Ride Problem with Synchronized Visits  
**Zhao, B., Puchinger, J., Leus, R.** ..... 240
- Fast Passenger Assignment for Public Transport Services  
**Wens, M., Verstraete, J., Vansteenwegen, P.** ..... 241
- Column elimination for scheduling problems  
**Coppé, V.** ..... 243
- Integrating Public Transportation in Bike Sharing Network Design: A Mathematical Model and Numerical Analysis  
**Bayri, S., Braekers, K., Caris, A.** ..... 245

*13h20–14h40* **Parallel session FC3**

*Session chair:* Dimitri Papadimitriou

*Room:* HOGM 00.85

ORBEL Award session

- De Latte, J., Reussens, Q., Marx, A., Goossens, L.** .... 247

*13h20–14h40* **Parallel session FC4**

*Session chair:* Yves Sagaert

*Room:* HOGM 01.85

- Integrating stochastic optimization and decision-focused learning  
**Barbini, A., De Baets, B., Verwaeren, J.** ..... 248

Temporal Model-Agnostic Meta-Learning for Cold-Start Time Series Forecasting

*Janssens, W., Bogaert, M., Van den Poel, D.* ..... 250

Degradation-aware spare parts management using Deep Reinforcement Learning

*Al Houry, N., Claeys, D., Bukhsh, Z.* ..... 252

Predictive Maintenance for Industrial Cinema Projectors Using Machine Learning

*Dam, S., Coussement, K.* ..... 254

**13h20–14h40 Parallel session FC5**

*Session chair:* Phablo Moura

*Room:* HOGS 02.10

Determining extremal chemical graphs of maximum degree at most 3 via a polyhedral description

*Dusollier, V., Bonte, S., Devillez, G., Hertz, A., Mélot, H., Schindl, D.* ..... 256

On the approximability of the balanced minimum evolution problem

*Catanzaro, D., Pesenti, R., Pisanu, F.* ..... 258

Genetic Branching: An Evolutionary Framework for Interpretable Branching Strategies

*Renard, S., Louveaux, Q., Fortz, B.* ..... 260

A New Algorithm to Enumerate the Vertices of the Balanced Minimum Evolution Polytope

*Pierre, B., Catanzaro, D., Ronco, R.* ..... 261

**15h10–16h10 Plenary talk**

*Room:* STUK 02.C004

**Stefan Feuerriegel**

Causal ML to optimize decisions from observational data

*Session chair:* Wouter Verbeke

**16h10–16h30 ORBEL Award announcement and closing session**

*Room:* STUK 02.C004

*Session chairs:* Dimitri Papadimitriou and Wouter Verbeke





# Abstracts

TA1 .....	39
TA2 .....	45
TA3 .....	53
TA4 .....	61
TA5 .....	67
TB1 .....	72
TB2 .....	80
TB3 .....	88
TB4 .....	94
TB5 .....	102
TC1 .....	108
TC2 .....	112
TC3 .....	118
TC4 .....	124
TC5 .....	130
TD1 .....	136
TD2 .....	142
TD3 .....	148
TD4 .....	149
TD5 .....	154
FA1 .....	158
FA2 .....	165
FA3 .....	171
FA4 .....	178
FA5 .....	186
FB1 .....	194
FB2 .....	201
FB3 .....	209
FB4 .....	217
FB5 .....	225
FC1 .....	233
FC2 .....	240
FC3 .....	247
FC4 .....	248
FC5 .....	256
<b>List of Authors .....</b>	<b>263</b>



# Computer Vision for Anti-Money Laundering in Time-Varying Transaction Networks

Bruno Deprez

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

University of Antwerp - imec, Department of Mathematics

e-mail: [bruno.deprez@kuleuven.be](mailto:bruno.deprez@kuleuven.be)

Tim Verdonck

University of Antwerp - imec, Department of Mathematics

KU Leuven, Department of Mathematics

e-mail: [tim.verdonck@uantwerpen.be](mailto:tim.verdonck@uantwerpen.be)

Wouter Verbeke

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [wouter.verbeke@kuleuven.be](mailto:wouter.verbeke@kuleuven.be)

Network analytics has become an indispensable tool for anti-money laundering (AML), given the complexity of criminal schemes involving numerous transactions and parties to conceal illicit funds. Over the years, the focus has shifted from analyzing individual network positions via centrality measures to capturing complex interactions using graph neural networks (GNNs) and flow detection methods (Deprez et al., 2025b).

A segment of the literature focuses on identifying specific network motifs associated with money laundering. Although GNNs can be extended to capture complex motifs and topologies (Egressy et al., 2024), many existing approaches are tailored to very specific patterns (Lee et al., 2020; Li et al., 2020; Deprez et al., 2025a). Furthermore, these methods often utilize minimal information, incorporating at most the transaction amount and direction.

In this work, we propose a novel method to represent sub-graphs as visual images to capture intricate patterns within the transaction network. We build upon the work of Deprez et al. (2025a), who demonstrated that effective, scalable AML can be achieved using information restricted to the second-order neighbourhood of a node. We combine various summary statistics by carefully defining distinct channels in our visual representation, moving beyond the standard RGB format. Specifically, we map the node’s local topology onto a fixed-size grid, where surrounding pixels encode the aggregated statistics of neighbours. By utilizing high-dimensional tensors, we explicitly link these statistics; for example, the minimal and maximal values of a specific neighbourhood sub-region share the same pixel coordinates but occupy different channels. This contrasts with classic tabular data, where such spatial relationships are lost among feature columns.

We address the dynamic nature of money laundering by generating temporal snapshots of the transaction network, resulting in a video-like sequence of trans-

action behaviour for each node. We introduce *transaction echoes*, where transactions are not seen as binary connections (merely present or absent). Instead, past transactions remain present but their signal strength decays exponentially over time. This decay mechanism acts as a continuous memory state, overcoming the limitations of rigid time-windowing approaches which often sever the link between related transactions occurring across window boundaries. This enables the analysis of network structures and transaction flows that span multiple time steps.

Our experiments combine different approaches to analyse these feature videos and compare them against several baseline methods. The results demonstrate that our novel approach yields high detection rates comparable to complex GNN architectures, but with significantly lower computational overhead. This suggests that summarising network structures based on simplified visual statistics offers a robust and efficient alternative for real-world AML systems.

## References

- Bruno Deprez, Bart Baesens, Tim Verdonck, and Wouter Verbeke. Garg-aml against smurfing: A scalable and interpretable graph-based framework for anti-money laundering. *arXiv preprint arXiv:2506.04292*, 2025a. URL <https://arxiv.org/abs/2506.04292>.
- Bruno Deprez, Toon Vanderschueren, Bart Baesens, Tim Verdonck, and Wouter Verbeke. Network analytics for anti-money laundering—a systematic literature review and experimental evaluation. *INFORMS Journal on Data Science*, pages 1–36, 2025b. doi: 10.1287/ijds.2024.0042. URL <https://pubsonline.informs.org/doi/abs/10.1287/ijds.2024.0042>.
- Béni Egressy, Luc von Niederhäusern, Jovan Blanuša, Erik Altman, Roger Wattenhofer, and Kubilay Atasu. Provably powerful graph neural networks for directed multigraphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(10):11838–11846, Mar. 2024. doi: 10.1609/aaai.v38i10.29069. URL <https://ojs.aaai.org/index.php/AAAI/article/view/29069>.
- Meng-Chieh Lee, Yue Zhao, Aluna Wang, Pierre Jinghong Liang, Leman Akoglu, Vincent S Tseng, and Christos Faloutsos. Autoaudit: Mining accounting and time-evolving graphs. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 950–956. IEEE, 2020. doi: 10.1109/BigData50022.2020.9378346.
- Xiangfeng Li, Shenghua Liu, Zifeng Li, Xiaotian Han, Chuan Shi, Bryan Hooi, He Huang, and Xueqi Cheng. Flowscope: Spotting money laundering based on graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 4731–4738, 2020. doi: 10.1609/aaai.v34i04.5906. URL <https://ojs.aaai.org/index.php/AAAI/article/view/5906>.

# Robust decision-making for budget allocation in marketing using machine learning

Daan Caljon

KU Leuven, Faculty of Economics & Business

e-mail: [daan.caljon@kuleuven.be](mailto:daan.caljon@kuleuven.be)

Abdurahman Maarouf

LMU Munich & Munich Center for Machine Learning

Jente Van Belle

KU Leuven, Faculty of Economics & Business

Wouter Verbeke

KU Leuven, Faculty of Economics & Business

Stefan Feuerriegel

LMU Munich & Munich Center for Machine Learning

Marketing budgeting, which involves the allocation of resources across campaigns, regions, product lines, or even individual customers, is among the most critical financial decisions firms must make (Piercy, 1987). In many industries, marketing spending accounts for a substantial share of total expenditures. In some sectors, such as consumer package goods, it even represents more than 25% of the total budget (The CMO Survey, 2025). This large share not only highlights the importance of marketing spending for driving sales and brand equity, but it also implies that inefficient budgeting can have detrimental financial consequences (Edeling and Fischer, 2016).

When estimating the impact of marketing budget allocation decisions on revenue, it is crucial to consider *incremental effects* (Ascarza, 2018). For example, some segments might have high revenue regardless of marketing budget spent (Silva-Risso et al., 1999). Consequently, the budget should not be allocated to these segments, but to those with the largest estimated incremental revenue effect as to maximize the total revenue increase. In practice, firms often use heuristics to allocate budgets (Gupta and Steenburgh, 2008). For example, many companies use the percent-of-sales rule (Sinha and Zoltners, 2001; Gupta and Steenburgh, 2008). This heuristic allocates budgets such that the ratio of budget spent compared to revenue for that segment is smaller than some set value. However, this approach ignores heterogeneity between segments. Importantly, capturing such heterogeneity is crucial when maximizing the gains from budget allocation, as some segments will benefit more from additional budgets than others.

Another important concern is that, for some segments, the point estimate of the incremental effect may be high, but the corresponding prediction interval may also be large, meaning that, in the worst case, firms run the risk of los-

ing a significant amount of money by allocating large budgets to such segments. Nevertheless, prior approaches for budget allocation do not account for this uncertainty. This is undesirable for several reasons. First, the costs and benefits of a marketing initiative might be asymmetric. For example, an advertisement might make potential customers not buy the product, while they would have bought it if the ad had not existed. Second, decision-makers are often risk-averse, meaning they are willing to trade off some expected value for less uncertainty.

To this end, we propose optimizing marketing budgets using a two-step approach. First, the heterogeneous incremental effects of marketing budget allocation are estimated using causal machine learning (Wager and Athey, 2018). Second, we employ robust optimization (Bertsimas and Sim, 2004) to allocate marketing budgets, specifically accounting for the uncertainty in the incremental effect estimates. We validate this approach using real-world data from a US retail chain.

## References

- Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of marketing Research*, 55(1):80–98.
- Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations research*, 52(1):35–53.
- Edeling, A. and Fischer, M. (2016). Marketing’s impact on firm value: Generalizations from a meta-analysis. *Journal of Marketing Research*, 53(4):515–534.
- Gupta, S. and Steenburgh, T. (2008). *Allocating marketing resources*. SSRN.
- Piercy, N. F. (1987). The marketing budgeting process: marketing management implications. *Journal of Marketing*, 51(4):45–59.
- Silva-Risso, J. M., Bucklin, R. E., and Morrison, D. G. (1999). A decision support system for planning manufacturers’ sales promotion calendars. *Marketing Science*, 18(3):274–300.
- Sinha, P. and Zoltners, A. A. (2001). Sales-force decision models: Insights from 25 years of implementation. *Interfaces*, 31(3\_supplement):S8–S44.
- The CMO Survey (2025). Marketing expenses as a percentage of companies’ overall budgets according to chief marketing officers (cmos) in the united states as of february 2025, by industry. <https://www.statista.com/statistics/742988/marketing-budget-share-category-usa/>. Graph. In Statista. Retrieved May 05, 2025.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

# Entropy-Based Model Evaluation for Guided Contextualized Process Discovery

Zahra Ahmadi

KU Leuven Research Centre for Information Systems Engineering (LIRIS)

e-mail: [zahra.ahmadi@kuleuven.be](mailto:zahra.ahmadi@kuleuven.be)

Jochen De Weerd

KU Leuven Research Centre for Information Systems Engineering (LIRIS)

e-mail: [jochen.deweerd@kuleuven.be](mailto:jochen.deweerd@kuleuven.be)

Estefanía Serral

KU Leuven Research Centre for Information Systems Engineering (LIRIS)

e-mail: [estefania.serralasensio@kuleuven.be](mailto:estefania.serralasensio@kuleuven.be)

Contextualized process discovery [1] enriches event logs with contextual information such as resources, location, or case attributes to obtain process models that better reflect behavioral variation [2]. A persistent challenge, however, is selecting an appropriate level of contextualization during model discovery. If too little context is introduced, meaningful differences between executions may remain hidden. If too much context is introduced, the resulting models can become overly complex and hard to interpret. As a result, analysts often rely on trial-and-error and expert intuition to navigate the trade-off between informativeness and interpretability [3].

We propose an entropy-based model evaluation approach that provides objective support for this selection problem. Our key contribution is the introduction of a pair of complementary, efficiently computable metrics, namely Directly Follows Graph (DFG) entropy and Entropic Relevance (ER). These metrics allow experts to compare alternative contextualized models of the same process and choose a specialization level aligned with their analytical goal. We provide a ranking and a Pareto-style comparison of candidate contextualizations. The first metric quantifies the structural complexity of the discovered model, while the second captures behavioral variability in the observed executions. Together, they position candidate models in a shared complexity–informativeness view and make the trade-offs explicit and reproducible.

The first algorithm computes DFG entropy on the Directly Follows Graph (DFG). It measures how predictable activity-to-activity transitions are by analyzing the diversity and distribution of incoming and outgoing directly-follows relations for each activity. High DFG entropy indicates many plausible next (or previous) steps with a more variable and branching structure. Low DFG entropy indicates more deterministic flows and clearer structure. This metric, therefore, reflects how contextualization affects structural interpretability.

The second algorithm computes Entropic Relevance (ER) [4], which captures

behavioral unpredictability at the log level. Unlike DFG entropy, ER is not limited to local transition choices. It reflects how diverse the observed executions are and how well they can be explained by the transition patterns implied by the model. Lower ER values correspond to more regular and predictable behavior, while higher values indicate substantial variability.

Given a set of contextualization or specialization options proposed by analysts or domain experts, we compute both DFG entropy and ER for each resulting candidate model. This yields a set of models that can be compared along two dimensions: structural complexity and behavioral variability. Experts can then select a model that balances interpretability and context-sensitivity according to the use case, while reducing manual effort and making the contextualization decision transparent.

**keywords:** Process Mining, contextualized process discovery, entropy, model evaluation, interpretability.

## References

- [1] Ahmadi, Z., De Weerd, J., and Serral, E.: A Novel Contextualization Method for Process Discovery Using Activity Specialization Hierarchies. *International Conference on Business Process Modeling, Development and Support* (2024).
- [2] Koschmider, A., Mannhardt, F., Heuser, T.: On the contextualization of event-activity mappings. In: *BPM 2018, Sydney, NSW, Australia, September 9-14, 2018, Revised Papers 16*. pp. 445–457. Springer (2019).
- [3] Ahmadi, Z., De Weerd, J., and Serral, E.: Context-driven process discovery: enhancing process flow interpretability with contextualized activity hierarchies. *Software and Systems Modeling* (2025).
- [4] Alkhamash, H., Polyvyanyy, A., Moffat, A., García-Bañuelos, L.: Entropic relevance: A mechanism for measuring stochastic process models discovered from event data. *Information Systems* 107, 101922 (2022).

# Generating Realistic Solar Irradiation Scenarios with Limited Historical Data

Robert Hanusa, Birger Raa, Stijn De Vuyst

Ghent University, Industrial Systems Engineering

`robert.hanusa@ugent.be`, `birger.raa@ugent.be`, `stijn.devuyst@ugent.be`

5 December, 2025

As companies increasingly look to integrate on-site production of variable renewable energy (VRE) in their production planning, it becomes important to use models that take into account the uncertainty of VRE forecasts. This research is inspired by the e-CODUCT project [1] and focuses on the case where an industrial user with a photovoltaic (PV) array wants to plan the energy usage for the coming day based on a point forecast. To address uncertainty in the forecast, multiple scenarios can be generated that reflect different possible realizations. We investigate a case in which the PV array is new, or where the company is in the planning phase and simulating performance, and therefore possesses very limited historical data for its specific location. While historical weather measurements may be available from nearby stations, the corresponding historical PV forecast data either does not exist or is only available in a spatially aggregated form. Aggregated PV forecasts behave differently from actual location-specific PV production, as small installations are more sensitive to localized events such as cloud passages.

Given this situation where there is limited or no historical forecast data, but a point forecast for the upcoming day is available, it is clear that several scenario generation approaches commonly used in other research [2] are not suitable:

- Deep Generative Models (DGMs): these methods require large volumes of training data and are therefore not appropriate for a data-scarce setting.
- Dependence-modelling via Copulas: Copulas normally require both historical forecasts and actuals. Without this, forecast errors cannot be estimated reliably, making standard copula-based approaches difficult to apply.
- Time-series generative models (e.g., ARIMA): These models generally generate future values using past observations and are therefore unsuitable for generating scenarios that reflect the next day's forecast.

Other studies face a similar constraint and therefore generate scenarios directly from a point forecast. These approaches typically rely on additive-noise models, such as sampling from parametric distributions (e.g., Beta or Gaussian) fitted to residuals. However, the resulting scenarios end up with a much higher point-to-point difference (average of the absolute differences between adjacent

points) than the original forecast, leading to scenarios that do not properly reflect the properties of the forecast well.

In this research, we propose a technique we call the Gamma Shift in which we use random numbers from the Gamma distribution to vary the forecast on the time axis rather than the magnitude axis. This has the benefits that (1) the point-to-point difference of the scenarios is similar to the original forecast, leading to more realistic scenarios and (2) the new scenarios reflect realistic cases where, for example, a cloud passes overhead at a different time than originally forecast, which is a very real possibility.

Figure 1 shows an example of the solar irradiance of three scenarios with a standard additive-noise technique (left) and with the Gamma Shift technique (right). The forecast (solid black line) is shown against three scenarios generated using additive noise (dashed lines) from a normal distribution with mean 0 and standard deviation  $50W/m^2$ . Visually, the forecast is a much smoother curve than the scenarios, which appear more erratic in shape. Similar results would be seen for the additive-noise technique with other probability distributions, as long as the average is zero and the samples are independent and identically distributed.

For the Gamma Shift scenarios on the right, we see that the three scenarios (dashed lines) are shifted in the time domain, leading to curves with a smoothness, as measured by point-to-point difference, similar to that of the forecast's (solid black line).

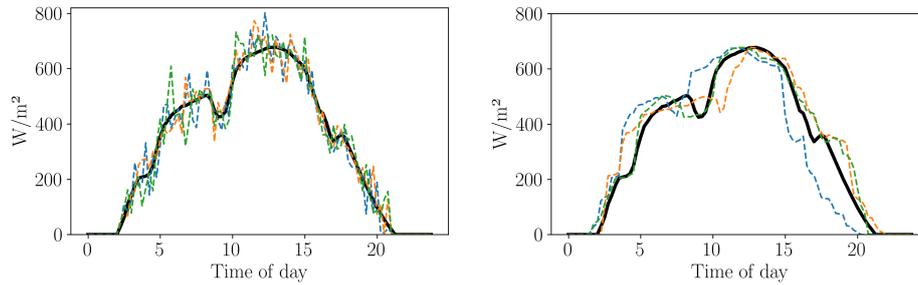


Figure 1: Scenarios generated with additive noise (left) and the Gamma Shift (right). Daylight hours extended from 2:00 to 22:00 for visual effect.



The e-CODUCT project is funded under Horizon Europe Grant Agreement n°101058100.

## References

- [1] European Commission (2025). *Horizon Europe Grant Agreement No. 101058100*. Retrieved from <https://e-coduct.eu/>.
- [2] Kousounadis-Knousen, M. A., Bazionis, I. K., Georgilaki, A. P., Catthoor, F., and Georgilakis, P. S. (2023). A review of solar power scenario generation methods with focus on weather classifications, temporal horizons, and deep generative models. *Energies*, 16(15), 5600. <https://www.mdpi.com/1996-1073/16/15/5600>. doi:10.3390/en16155600.

# Sensitivity analysis for linear changes of the constraint matrix of a (mixed-integer) linear program

Bardhyl Miftari

University of Liège, Montefiore Institute

e-mail: [bmiftari@uliege.be](mailto:bmiftari@uliege.be)

Damien Ernst

University of Liège, Montefiore Institute

e-mail: [dernst@uliege.be](mailto:dernst@uliege.be)

Quentin Louveaux

University of Liège, Montefiore Institute

e-mail: [q.louveaux@uliege.be](mailto:q.louveaux@uliege.be)

Guillaume Derval

University of Liège, Montefiore Institute

e-mail: [gderval@uliege.be](mailto:gderval@uliege.be)

December 2025

One strength of linear programming models is the ability to analyze the variation of an optimal solution when we account for slight changes in the data. A popular example comes from modifying one cost coefficient in the objective. In this case, sensitivity analysis allows you to determine in which range the cost coefficient can vary while keeping the current optimal solution optimal. Similarly, if you assume the variation of one right-hand-side value, we can compute the interval in which the current optimal basis remains optimal. In this case, though, the optimal value varies linearly and the optimal dual variable provides the slope of such a linear variation. This can be very helpful for practitioners as it helps understand how robust their solution is and identify which parameters have the greatest impact on the outcome. As a result, linear programming not only provides an optimal solution but also valuable insights into how changes in the model's inputs affect that solution.

However the analysis becomes more complicated when we try to extend it to changing several coefficients simultaneously, or when we want to analyze the variation implied by the change of a matrix coefficient, especially when this matrix coefficient appears in the basis. In this talk, we propose a new model for studying the variation of the optimal value on a linear program. We assume that one parameter varies, but this single parameter may appear in several coefficients of

the matrix  $A$ . More specifically, we consider a linear (or mixed-integer) program

$$\begin{aligned} f(\lambda) = \min \quad & \mathbf{c}^t \mathbf{x} \\ \text{s.t} \quad & A\mathbf{x} + \lambda D' \mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \in \mathbb{R}^{n_c} \times \mathbb{Z}^{n_i}. \end{aligned} \tag{1}$$

where  $\lambda \in [\lambda_1, \lambda_2]$  models the uncertainty, more specifically the potential variation of the coefficients within some limits.

Given the hardness of computing  $f(\lambda)$  exactly, we propose methods to compute upper and lower bounds on  $f(\lambda)$ . The bounding methods are based on two approaches: robust optimisation and Lagrangian relaxations. Out of each of these two approaches, we derive bounding methods of three types: ones that compute constant bounds, ones giving variables bounds depending on  $\lambda$  and ones computing envelopes on the objective function. The envelopes are a combination of multiple variable bounds, proven optimal in the sense that no other variable bound of the same type is tighter, that form an envelope around the objective function. Some methods add new degrees of freedom, and therefore, can lead to multiple variations. We also introduce an iterative algorithm to compute these bounds with relatively small gaps and discuss its efficiency.

We test our methods on a newly created benchmark of instances inspired by real-life problems. The computational results show that we are able, in most cases, to obtain accurate bounds around  $f(\lambda)$  in much less time than what is needed to sample 100 optimal points, while providing a useful guarantee that cannot be obtained from the sampling.

# A Distance Metric for Mixed Integer Programming Instances

Gwen Maudet

SnT, University of Luxembourg, Esch-sur-Alzette, Luxembourg

Grégoire Danoy

FSTM/DCS, SnT, University of Luxembourg, Esch-sur-Alzette, Luxembourg

A solving strategy (SS) within a Mixed-Integer Linear Programming (MILP) solver can be defined by components that guide decisions during the solving process until optimality is reached. With the rapid growth of machine learning (ML), an SS can be formalized as a classification problem and learned automatically. However, existing models, while effective on homogeneous and predefined instance sets, struggle to generalize to more diverse or structurally different instances. A significant gap therefore remains before ML-based SSs can be reliably integrated into general-purpose solvers.

Introducing a distance metric between MILP instances is a key step toward this generalization. Such a metric quantifies the heterogeneity of instance sets used to evaluate SSs, providing a principled indicator of a model’s ability to generalize across the MILP space. On the other side, current approaches perform well on homogeneous groups but degrade when considering the full diversity of MILP instances. Segmenting the MILP space into homogeneous regions using a suitable distance would then offer one way to have a generalizable framework for SS using ML, by assigning a dedicated SS to each region.

Several methods have been proposed to measure similarity between MILP instances. The *ISS* approach [2] encodes constraint matrices as images analyzed by autoencoders, but is sensitive to constraint ordering. The *GNV* approach [3] provides an invariant representation and improves generalization within predefined classes, but requires supervised training and remains limited to those classes. The MIPLIB 2017 *Feat* distance [1] relies on more than 100 handcrafted features, yet lacks theoretical grounding for feature selection or normalization.

To overcome these limitations, we propose a distance, *Formal*, directly derived from the mathematical formulation of MILP instances, requiring no supervised training; this extended abstract is a condensed version of [4]. Variables, coefficients, and right-hand-side terms are categorized, enabling a constraint representation based on the proportions of variable–coefficient pairs, and an instance representation based on the proportions of similar constraints. This yields a size-independent representation suitable for comparing instances of different dimensions. We quantify distances between constraints using the earth mover’s distance [5], and extend this principle to entire instances.

Our experiments assess the ability of the metric to identify as similar the instances belonging to the same class, across five classes table 1. The full study [4] includes 19 classes and additional subclass analyses. We evaluate multiple con-

figurations: ablations ignoring weight ( $\alpha$ ), variable ( $\beta$ ), right-hand-side ( $\gamma$ ), or objective ( $\zeta$ ) components, as well as exact ( $\emptyset E$ ) and approximate ( $\emptyset$ ) variants. Results show that considering all components of the formulation is crucial for accurate similarity estimation. The approximate version is on average 200 times faster while producing results comparable to the exact version, and it exhibits fewer time variations. Compared to the unsupervised baselines *ISS* and *Feat*, our approach significantly outperforms them and matches the performance of the supervised *GNN* specifically trained on the considered classes.

	Formal						Feat	ISS	GNN
	$\alpha$	$\beta$	$\gamma$	$\zeta$	$\emptyset E$	$\emptyset$			
bpp	42	47	49	26	47	47	27	32	81
bp2	62	71	100	71	95	95	24	52	100
bif	50	44	50	50	50	50	21	32	43
clp	54	59	37	64	57	56	28	12	37
col	40	40	40	40	40	40	27	36	40

Table 1: Evaluation of different similarity measures.

In conclusion, we introduce a formal MILP distance that serves as a foundational tool for instance analysis and integration into ML pipelines for SSs. Future work includes developing a portfolio scheme in which groups of instances are associated with ML-designed SSs, enabling automatic selection of the most suitable SS for any new instance.

**Ack.:** This research was funded by ANR France (ANR-22-CE46-0011) and FNR Luxembourg (INTER/ANR/22/17133848) through the UltraBO project.

## References

- [1] Gleixner, Ambros and Hendel, Gregor and Gamrath, Gerald and Achterberg, Tobias and Bastubbe, Michael and Berthold, Timo and Christophel, Philipp and Jarck, Kati and Koch, Thorsten and Linderoth, Jeff and Lübbecke, Marco and Mittelmann, Hans D. and Ozyurt, Derya and Ralphs, Ted K. and Salvagnin, Domenico and Shinano, Yuji: MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation* (2017).
- [2] Steever, Zachary and Murray, Chase and Yuan, Junsong and Karwan, Mark and Lübbecke, Marco: An Image-Based Approach to Detecting Structural Similarity Among Mixed Integer Programs. *INFORMS Journal on Computing* (2022).
- [3] Steever, Zachary and Hunt, Kyle and Karwan, Mark and Yuan, Junsong and Murray, Chase C.: A Graph-Based Approach for Relating Integer Programs. *INFORMS Journal on Computing* (2024).
- [4] Maudet, Gwen and Danoy, Grégoire: A Distance Metric for Mixed Integer Programming Instances. *European Conference on Artificial Intelligence* (2025).
- [5] Rubner, Yossi and Tomasi, Carlo and Guibas, Leonidas J: The Earth Mover’s Distance as a Metric for Image Retrieval. *International Journal of Computer Vision* (2000).

# Approximation algorithms for graph search problems with imperfect detection

Martijn van Ee

Netherlands Defence Academy, Faculty of Military Sciences

e-mail: [M.v.Ee.01@mindef.nl](mailto:M.v.Ee.01@mindef.nl)

René Sitters

Vrije Universiteit Amsterdam, Department of Operations Analytics

e-mail: [r.a.sitters@vu.nl](mailto:r.a.sitters@vu.nl)

Imperfect detection is a common feature in search theory [2], and is generally motivated by inaccuracies in sensor measurements or the inability to perfectly sweep an area. In search problems, one is usually given a set of cells, and the goal is to design a search strategy that minimizes the expected time until the target has been found. Alternatively, a common objective is the probability of finding the target, given a certain deadline. However, travel times are usually ignored in this field. On the other hand, searching with travel times is well studied in graph search problems, such as the traveling repairman problem (TRP) [3] and the expanding search problem [5]. In these problems, detection is assumed to be perfect. In this article, we combine both fields, with a focus on approximation algorithms. In particular, we extend the approximability results for graph search problems with perfect detection to the case with imperfect detection. These results were published in [4].

In our models, we assume the searcher moves at unit speed through a weighted graph, and that searching a vertex takes 1 time unit. The probability that the target is detected while it is at the vertex that is being searched is  $\gamma$ . We consider the case where the target is hidden at random, and the case where the target is hidden adversarially. If the target is hidden at random, our goal is to find an (infinite) search path that minimizes the expected time until the target is found. If the target is hidden adversarially, meaning that it can observe our search strategy before choosing a vertex to hide, our goal is to find an (infinite) search path that minimizes the maximum expected time until detection over the vertices.

We consider two ways to search the graph: *pathwise search* (as is standard in routing problems) and *expanding search*, where we only incur the cost of an edge if we use it for the first time and any following traversal has no cost. Natural applications of searching with imperfect detection include search and rescue missions, where a local search (such as an excavation) as well as moving equipment takes substantial time. Other applications can be found in testing the state of a system by investigating its components. The reliability of a test depends on the thoroughness and precision and hence on the time invested. Moreover, switching between tests takes a setup time. This is closely related to time-critical testing problems [1].

**Our contributions** Our first result is that the problem with random targets and pathwise search reduces, with a loss of a factor  $1 + \varepsilon$ , to weighted-TRP. Since there exists a  $(3.59 + \varepsilon)$ -approximation for weighted-TRP, we obtain an  $(3.59 + \varepsilon)$ -approximation. It turns out that roughly the same reduction also works for expanding search, by which we obtain an  $(2e + \varepsilon)$ -approximation. As a side result, we note that a similar reduction to the orienteering problem also gives the first constant factor approximation for the problem of maximizing the probability of finding the target, given a certain deadline.

For imperfect graph search problems with adversarial targets, we present simple algorithms with constant factor approximation guarantees. For pathwise search, the algorithm’s solution walks back and forth along the (approximate) shortest path that visits all vertices, and either searches for one time unit or for  $\lceil 1/\gamma \rceil$  time units when visiting a vertex. This gives a 3.2826-approximation. For expanding search, we construct the minimum spanning tree in an expanding manner, and search each vertex for one time unit. Afterwards, the solution searches the vertices in reversed order, and so on. This gives a 2.0876-approximation. In Table 1 is a summary of our results, together with known approximability results for perfect detection.

Table 1: Approximation ratios for imperfect detection compared to approximation ratios for perfect detection.

	Random target		Adversarial target	
	Perfect	Imperfect	Perfect	Imperfect
Pathwise search	$3.59 + \varepsilon$	$3.59 + \varepsilon$	$1.5 - \varepsilon_0$	3.2826
Expanding search	$2e + \varepsilon$	$2e + \varepsilon$	1	2.0876
Orienteering	$0.5 - \varepsilon$	$0.5 - \varepsilon$		

## References

- [1] Alessandro Agnetis, Ben Hermans, Roel Leus, and Salim Rostami. Time-critical testing and search problems. *European Journal of Operational Research*, 296(2):440–452, 2022.
- [2] Stanley J. Benkoski, Michael G. Monticino, and James R. Weisinger. A survey of the search theory literature. *Naval Research Logistics*, 38(4):469–494, 1991.
- [3] Kamalika Chaudhuri, Brighten Godfrey, Satish Rao, and Kunal Talwar. Paths, trees, and minimum latency tours. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 36–45, 2003.
- [4] Martijn van Ee and René Sitters. Approximation algorithms for graph search problems with imperfect detection. In *Proceedings of the 23rd International Workshop on Approximation and Online Algorithms*, pages 213–227, 2025.
- [5] Svenja M. Griesbach, Felix Hommelsheim, Max Klimm, and Kevin Schewior. Improved approximation algorithms for the expanding search problem. In *Proceedings of the 31st Annual European Symposium on Algorithms*, pages 54:1–54:15, 2023.

# DDOLib: A library for Solving Discrete Optimization Problems with Dynamic Programming

Pierre Schaus<sup>1</sup>, Quentin Meurisse<sup>2</sup>, Roger Kameugner<sup>1</sup>,  
Renaud De Landtsheer<sup>2</sup>, Fabian Germeau<sup>2</sup>, Xavier Gillard<sup>1</sup>

<sup>1</sup>UCLouvain, ICTeam, Belgium

<sup>2</sup>CETIC, COAL, Belgium

Emails:

`first.last@uclouvain.be` (UCLouvain authors),

`first.last@cetic.be` (CETIC authors)

DDOLib [1] is a Java library for solving discrete minimization problems using *Dynamic Programming* [2]. By implementing a model, the user specifies the state representation, the transition function, and the cost structure, and the solver automatically orchestrates the search. A library for modeling in this paradigm was introduced in the DDO solver [8], in DDIP solver [10], or CODD [11] solvers, but DDOLib has the ambition to unify them all.

Conceptually, solving the dynamic program is equivalent to finding the shortest path in a layered acyclic graph, from an initial state to a designated terminal state. A broad range of algorithms has been developed for discovering shortest path problems while avoiding to expand explicitly the whole search space. DDOLib supports several such informed search strategies for this purpose, including A\* [9], but also any-time-complete algorithm such as the *Anytime Column Search* [12] or *Anytime Weighted A\** [13]. These approaches become particularly efficient when the user provides an admissible or consistent heuristic that offers an optimistic estimate of the remaining cost to reach the target state.

Beyond classical search-based dynamic programming, DDOLib integrates decision diagram based solving techniques [3, 4]. By exploring the state space with branch-and-bound, this approach interleaves two bounded width decision diagrams. The first is the *restricted decision diagram* obtained by discarding less promising states that provide a primal bound. The second is the *relaxed decision diagram* obtained by state merging that yields a dual bound. The latter requires, in addition to the transition function, that the user defines a state-merge operator to combine nodes during diagram construction. In addition, DDOLib incorporates all recent advances associated with these decision diagram based methods, including caching mechanisms [5], dominance detection [6], and local-bounds computations [7], among others.

In this presentation, we illustrate how to use DDOLib through several examples, including the well-known binary knapsack problem, various single-machine scheduling problems, and the traveling salesman problem with time windows.

**Acknowledgment:** This work is supported by the Win4Collective DDOLib project of the Walloon Region (project 2410118).

## References

- [1] DDOLib website. <https://github.com/DDOLIB-CETIC-UCL/DDOLib>.
- [2] Richard Bellman. The theory of dynamic programming. *Bull. Amer. Math. Soc.*, 60(6):503–515, 1954.
- [3] David Bergman, Andre A. Cire, Willem-Jan Van Hoeve, and John Hooker. *Decision Diagrams for Optimization*. Springer, 2016.
- [4] David Bergman, Andre A. Cire, Willem-Jan Van Hoeve, and John N. Hooker. Discrete optimization with decision diagrams. *IJOC*, 28(1):47–66, 2016.
- [5] Vianney Coppé, Xavier Gillard, and Pierre Schaus. Decision diagram-based branch-and-bound with caching for dominance and suboptimality detection. *IJOC*, 36(6):1522–1542, 2024.
- [6] Vianney Coppé, Xavier Gillard, and Pierre Schaus. Modeling and exploiting dominance rules for discrete optimization with decision diagrams. In *CPAIOR*, 2024.
- [7] Xavier Gillard, Vianney Coppé, Pierre Schaus, and André A. Cire. Improving the filtering of branch-and-bound mdd solver. In *CPAIOR*, 2021.
- [8] Xavier Gillard, Pierre Schaus, and Vianney Coppé. Ddo: A generic and efficient framework for mdd-based optimization. In *IJCAI*, pages 5243–5245, 2021.
- [9] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybern.*, 4(2):100–107, 1968.
- [10] Ryo Kuroiwa and J. Christopher Beck. Domain-independent dynamic programming: Generic state space search for combinatorial optimization. In *ICAPS*, volume 33, 2023.
- [11] Laurent Michel and Willem-Jan van Hoeve. Codd: A decision diagram-based solver for combinatorial optimization. *ECAI*, 2024.
- [12] Satya G. Vadlamudi, Piyush Gaurav, Sandip Aine, and P. P. Chakrabarti. Anytime column search. In *AJCAI*, 2012.
- [13] Rong Zhou and Eric A. Hansen. Multiple sequence alignment using anytime a\*. In *AAAI*, pages 975–977, 2002.

# Sequence Variables: A Constraint Programming Computational Domain for Routing and Sequencing – Extended Abstract

Augustin Delecluse

KU Leuven, CODeS

e-mail: [augustin.delecluse@kuleuven.be](mailto:augustin.delecluse@kuleuven.be)

Pierre Schaus

UCLouvain, ICTEAM

e-mail: [pierre.schaus@uclouvain.be](mailto:pierre.schaus@uclouvain.be)

Pascal Van Hentenryck

GATECH

e-mail: [pascal.vanhentenryck@isye.gatech.edu](mailto:pascal.vanhentenryck@isye.gatech.edu)

*This is an extended abstract of a paper still under review. A preprint version of the paper can be found in [1].*

## Abstract

Constraint Programming (CP) offers an intuitive, declarative framework for modeling Vehicle Routing Problems (VRP), yet classical CP models based on successor variables cannot always deal with optional visits or insertion based heuristics. We introduce a new computation domain in CP, that addresses these limitations called *insertion-based sequence variables* [1] and illustrate their use on a pickup and delivery problem as well as some comparative results.

## 1 Introduction

The pickup and delivery problem consists of scheduling a fleet of  $K$  vehicles to fulfill a set of transportation requests  $R$ , where each request specifies a pickup location and a delivery location. The objective is to minimize the total distance traveled.

This problem can be modeled with one insertion-based sequence variable for each vehicle. Each variable represents the sequence of nodes visited by the corresponding vehicle, starting at the depot and ending at the depot. Constraints are added to those variables to enforce restrictions on the paths, such as distance limit or vehicle capacity. Given the declarative nature of CP, a new constraint can be added within a few lines of code to an existing model, allowing to represent easily a new problem variant.

We give some insight on the filtering algorithms of some of the important constraints computational domain in charge of removing infeasible insertions in the sequences.

We also show how sequence variables can be instrumental at defining first-fail insertion-based backtracking depth-first-search strategies inspired by [2]. This process is illustrated in Figure 1, where a partial path (represented by a sequence variable) visiting request  $r_1$  is expanded to add a new request  $r_2$ .

Using these variables leads to improved performance compared to other CP models across several pickup and delivery problem variants, and they demonstrate favorable performance compared to off-the-shelf solvers such as OR-Tools and Hexaly.

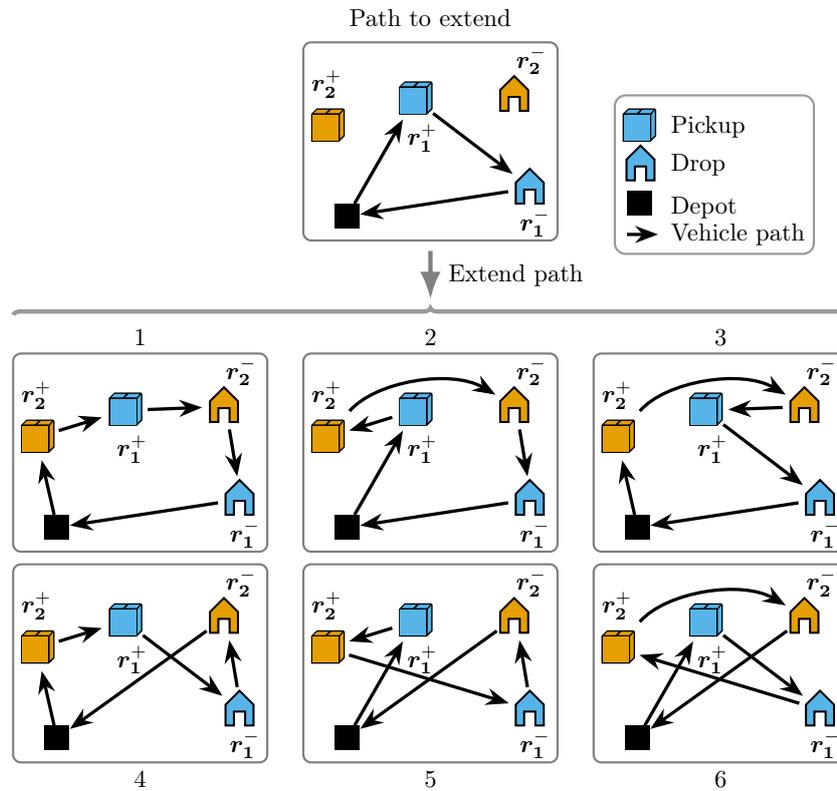


Figure 1: Paths generated from an initial one (top). Numbers indicate which are considered first.

## References

- [1] Augustin Delecluse, Pierre Schaus, and Pascal Van Hentenryck. Sequence variables: A constraint programming computational domain for routing and sequencing. *arXiv preprint arXiv:2510.09373*, 2025.
- [2] Philip Kilby, Patrick Prosser, and Paul Shaw. A comparison of traditional and constraint-based heuristic methods on vehicle routing problems with side constraints. *Constraints*, 5(4):389–414, 2000.

# Integrated location, partitioning and sequencing decisions for efficient power system restoration

Hatice Çalık

KU Leuven, ELECTA, Department of Electrical Engineering, Leuven  
ETCH, EnergyVille, Genk  
e-mail: [hatice.calik@kuleuven.be](mailto:hatice.calik@kuleuven.be)

Dirk Van Hertem

KU Leuven, ELECTA, Department of Electrical Engineering, Leuven  
ETCH, EnergyVille, Genk

Hande Yaman

KU Leuven, ORSTAT, Faculty of Economics and Business, Leuven

Power systems are increasingly challenged by environmental, technological, political, and socio-economic changes. Recent large blackouts in the Iberian Peninsula, Chile, Heathrow, and Berlin have shown that widespread outages can occur suddenly and cause major societal and operational disruption. Climate change is leading to more extreme weather events such as hurricanes, floods, and wildfires, which in turn increase the frequency of power outages [2]. Long-lasting blackouts create significant economic losses and serious societal impacts. The Iberian blackout of April 2025 (Spain and Portugal), reportedly caused by over-voltage [4], required 16 hours to restore the Spanish transmission system, while some areas remained without electricity for up to 24 hours. The event stranded around 35,000 railway and metro passengers, caused over 500 flight cancellations, and left many people without water or communication services. Economic losses are estimated between 1 and 4.5 billion EUR for Spain and around 2 billion EUR for Portugal. A similar blackout in Belgium would lead to an estimated loss of 2.7 billion EUR per working day.

These events highlight the importance of more resilient power systems with effective restoration procedures to reduce the impact of partial or total blackouts. Furthermore, fast and adaptable decision-making tools become essential to enable timely responses under damage and renewable uncertainties. To this end, our study investigates exact and efficient optimization methods for solving an NP-Hard power system restoration problem that consists of integrated facility location, network partitioning and sequencing decisions. The problem is a generalization of the parallel parallel power system restoration problem (PPSR) [3, 5] and the generator start-up sequencing (GSS) problem [1]. The GSS aims minimizing the latest startup time among all components in a power network after a complete blackout. The network consists of generators, critical loads (CLs) such as hospitals or military facilities as well as transshipment nodes. The generators are of two types: black start (BS) and non-black start (NBS), where only BSs can

restart themselves without demanding external power. The NBS and CL units can only start if they are provided with sufficient power. Moreover, the NBS units, such as Nuclear reactors, require the external power for a certain amount of time before they can actually start producing power themselves at a certain rate and feeding other network components. Thus, finding an optimal sequence requires choosing among a combinatorial set of options.

For technical and operational reasons, the power network is often divided into partitions which are referred to as *islands* before re-starting the components. The PPSR addresses the combined islanding and GSS on a network where the location of BSs are predetermined. Currently, the emerging storage technologies, particularly the battery electric storage systems (BESS) with large storage and BS capabilities, are increasingly considered as promising solutions for efficient and effective restoration towards the net-zero European grid of 2050. In the integration of BESS into the restoration procedures, determining their locations play an important role in minimizing the overall restoration time. The BESS location decisions combined with the PPSR constraints lead to a very challenging optimization problem. We investigate two variants of this integrated problem and develop integer programming-based solution methods. The computational analysis indicates that the overall restoration time can be improved with the optimal integration of BESS while also potentially phasing out a set of conventional BS units.

## References

- [1] Çalık, Hatice, Hakan Ergun, Kaan Yurtseven and Dirk Van Hertem. A bi-objective approach to power system restoration with renewable participation. IEEE 2024 18th International Conference on Probabilistic Methods Applied to Power Systems (PMAPS 2024), Date: 24-26 Jun 2024, Location: Auckland University of Technology, Auckland, New Zealand, 2024.
- [2] Climate Central, 2024. Weather-related Power Outages Rising. Accessed on: 23-01-2026. <https://www.climatecentral.org/climate-matters/weather-related-power-outages-rising>
- [3] Chopra, Sunil, Feng Qiu and Sangho Shim. Parallel Power System Restoration. INFORMS Journal on Computing, 35 (1), 233–247, 2023.
- [4] ENTSO-E, 2025. 28 April Blackout in Spain and Portugal: Expert Panel releases comprehensive factual report. Accessed on: 23-01-2026. <https://www.entsoe.eu/news/2025/10/03/28-april-blackout-in-spain-and-portugal-expert-panel-releases-comprehensive-factual-report/>
- [5] Qiu, Feng and Peijie Li. An integrated approach for power system restoration planning. Proceedings of the IEEE, 105 (7), 1234–1252, 2017.

# Arc-flow formulation for a two-stage two-dimensional dual bin packing problem for wood reuse

Pauline Bessemans, Célia Paquay, and Morgane Dumont

University of Liège, HEC Liège Management School

QuantOM, Research Centre for Quantitative Methods and Operations Management

e-mails: [pauline.bessemans@uliege.be](mailto:pauline.bessemans@uliege.be), [cpaquay@uliege.be](mailto:cpaquay@uliege.be), [morgane.dumont@uliege.be](mailto:morgane.dumont@uliege.be)

The increasing demand for raw materials such as wood undoubtedly contributes to the depletion of natural resources and global warming. To curb this phenomenon, a more sustainable and circular management of wood could be developed through the more efficient management of wood waste. This wood waste can be in the form of beams or pallets and could be considered as wooden slats. These slats could be combined, assembled, and glued to build Cross-Laminated Timber (CLT) panels for the construction industry, as shown in Figure 1.

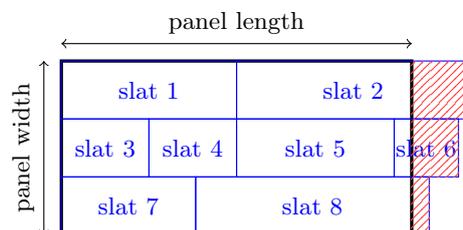


Figure 1: Toy example of assembly scheme to build one CLT panel

We aim to develop optimization techniques to recycle raw wood waste to create two-dimensional CLT panels of variable dimensions. The input waste is in the form of slats with variable and heterogeneous lengths and widths, but constant thickness. After a review of the Cutting & Packing literature, we identify our problem as an exact case of the two-stage two-dimensional multiple bin size dual bin packing problem (E-2S-2D-MBSDBPP) [1, 2] in the context of wood reuse optimization.

We propose a description of the problem and an arc-flow formulation with cuts, adapted from [5], [3], and [4]. The algorithm provides the assembly schemes that maximize the amount of wood reused through the production of CLT panels. Finally, we present the results of several numerical experiments based on realistic instances from the wood industry obtained by using Gurobi's B&B procedure.

Keywords: Dual bin packing, Cutting and Packing, Arc-flow formulation, Combinatorial Optimization

## References

- [1] Assmann, S. F., Johnson, D. S., Kleitman, D. J., & Leung, J.-T. (1984). On a dual version of the one-dimensional bin packing problem. *Journal of Algorithms*, 5(4), 502–525.
- [2] Gilmore, P. C. & Gomory, R. E. (1965). Multistage cutting stock problems of two and more dimensions. *Operations research*, 13(1), 94–120.
- [3] Martinovic, J., Delorme, M., Iori, M., Scheithauer, G., & Strasdat, N. (2020). Improved flow-based formulations for the skiving stock problem. *Computers & Operations Research*, 113, 104770.
- [4] Mrad, M. (2015). An arc flow-based optimization approach for the two-stage guillotine strip cutting problem. *Journal of the Operational Research Society*, 66(11), 1850–1859
- [5] Valério de Carvalho, J. (1999). Exact solution of bin-packing problems using column generation and branch-and-bound. *Annals of Operations Research*, 86(0), 629–659

## Multi-Commodity Inventory Routing Problem with Pickup and Delivery

Lukas Seeholzer

KU Leuven, Institute for Mobility - CIB

e-mail: [lukas.seeholzer@kuleuven.be](mailto:lukas.seeholzer@kuleuven.be)

Johan W. Joubert

KU Leuven, Institute for Mobility - CIB

e-mail: [johan.joubert@kuleuven.be](mailto:johan.joubert@kuleuven.be)

Kris Braekers

Hasselt University, Research Group Logistics

e-mail: [kris.braekers@uhasselt.be](mailto:kris.braekers@uhasselt.be)

Pieter Vansteenwegen

KU Leuven, Institute for Mobility - CIB

e-mail: [pieter.vansteenwegen@kuleuven.be](mailto:pieter.vansteenwegen@kuleuven.be)

The hospitality sector faces unique and challenging problems in logistics. Often, products have to be delivered under various circumstances. For example, a variety of different products from different suppliers needs to be delivered to a set of clients on a regular basis.

Based on a case study, a new variant of the inventory routing problem (IRP) is defined, modelled and analysed. We propose a mathematical model, show its limitations and propose a large neighbourhood search to find good solutions in a reasonable time, even for large instances.

The typical IRP combines the vehicle routing problem (VRP) with the inventory management on the client's side over multiple periods. Adding the time dimension to the search space massively increases the optimization opportunities compared to the VRP.

Most of the literature on variations of the IRP only considers one product to be delivered from one supplier[1]. Some works consider multiple suppliers[2] or multiple commodities[3], but very few consider both. In reality, however, there are often multiple suppliers that deliver different products to a set of clients, especially in the hospitality industry. Generally, there is no cooperation between the suppliers, leading to inefficiencies as each supplier operates its own vehicle fleet. Our case study considers a large hospitality service provider operating multiple production sites. In this setting, the clients are different production sites of the same company, while the suppliers are independent and deliver exactly one (type of) product to the customers each. This places this company in the unique situation of having knowledge about the operations of multiple suppliers and customers as well as the ability to take over the responsibility of organizing

the transportation.

This new situation not only allows to share vehicles among different suppliers, but adds the option to combine deliveries of different types of goods in one trip, including pickup and delivery. This even allows to deliver a total volume that is larger than the vehicle capacity during a single route, by picking up products from other suppliers after earlier deliveries. While the potential increase in efficiency is large, this freedom comes with significant optimization challenges since the route sequencing itself can influence the maximal amounts of the deliveries. To cover the aforementioned aspects, we define the multi-commodity IRP with pickup and delivery in the following way.

There is a set of customers and a set of suppliers. Each customer demands products from a subset of the suppliers. This demand may vary from day to day. Each supplier provides a specific (type of) product, which cannot be replaced by any other supplier. The clients have the possibility to store goods from any supplier. However, storage incurs costs. The goal is to choose a minimal vehicle fleet, operated centrally, and to minimize total costs, consisting of routing and storage costs. The optimization process needs to take decisions regarding routing, delivery schedules and storage while respecting constraints regarding demand satisfaction, vehicle capacities and route timing.

The combination of multiple time periods, multiple commodities provided by specific suppliers and inventory management leads to a novel problem definition that has not been studied yet, to the best of our knowledge.

Solving the problem with a mathematical model shows a significant cost saving potential due to an increase in vehicle efficiency. However, experiments show that finding the optimal solution, or even finding a feasible solution, becomes practically infeasible very quickly.

The high degree of freedom, spanning the deliveries, route sequencing and inventory management, create a huge solution space. Moreover, the mutual interdependence of many of the decisions that have to be taken at the same time lead to highly isolated feasible points.

To tackle these challenges, we present a metaheuristic to find good solutions in a reasonable time. The metaheuristic implements a large neighbourhood search with a variable neighbourhood descent and a destroy and repair move. We present the different moves constructed to cover the different optimization dimensions. With this approach, we achieve near optimal solutions for many instances and all obtained solutions provide large opportunities for optimization with respect to the current, uncooperative situation.

## References

- [1] Claudia Archetti, Nicola Bianchessi, Stefan Irnich, and M. Grazia Speranza. Formulations for an inventory routing problem. *International Transactions in Operational Research*, 21:353–374, May 2014.
- [2] Claudia Archetti, Marielle Christiansen, and M. Grazia Speranza. Inventory routing with pickups and deliveries. *European Journal of Operational Research*, 268(1):314–324, July 2018.
- [3] Jean-François Cordeau, Demetrio Lagana, Roberto Musmanno, and Francesca Vocaturo. A decomposition-based heuristic for the multiple-product inventory-routing problem. *Computers & Operations Research*, 55:153–166, March 2015.

## The optimization of store and warehouse based e-grocery fulfillment

Ruben D’Haen

Hasselt University, Research Group Logistics

e-mail: [ruben.dhaen@uhasselt.be](mailto:ruben.dhaen@uhasselt.be)

Charlotte Köhler

European University Viadrina, Department of Data Science and Decision Support

e-mail: [koehler@europa-uni.de](mailto:koehler@europa-uni.de)

Ann Melissa Campbell

University of Iowa, Department of Business Analytics

e-mail: [ann-campbell@uiowa.edu](mailto:ann-campbell@uiowa.edu)

Jan Fabian Ehmke

University of Vienna, Department of Business Decisions and Analytics

e-mail: [jan.ehmke@univie.ac.at](mailto:jan.ehmke@univie.ac.at)

After widespread adoption of e-commerce in the retail industry, online grocery sales by e-grocers are getting more popular as well. These e-grocers are struggling to optimize their fulfillment operations. In practice, two main modes of fulfillment can be identified: store and warehouse based fulfillment. In both modes, the requested products are first picked, i.e., collected from their storage locations, and afterwards delivered to the customer. In store based fulfillment, the existing network of grocery stores is used to pick the items, while in warehouse based fulfillment a dedicated warehouse is used for the picking operations.

Both store and warehouse picking have benefits and downsides. Grocery stores are usually located in the middle of the city, allowing for fast delivery after picking. However, order picking in stores is less efficient than in a warehouse, due to the layout of the store and the interactions with physical customers. While the warehouse allows for very efficient picking, it is usually located out of the city’s center. As such, the average delivery distance to customers is relatively large.

In our research, this trade-off between efficient picking and delivery is studied. We consider an integrated problem in which customer orders should be assigned to one of the fulfillment locations and the delivery routes, starting from an order’s assigned location, should be constructed. First, a mathematical model is used to optimize the operations on small problem instances. Next, a metaheuristic algorithm is developed and used to solve larger problem instances. To obtain practically relevant insights, a real-life data set is used to generate the e-grocery orders. Moreover, the algorithm is used to optimize the fulfillment operations in different problem settings to identify the critical parameters influencing the order assignment decision.

## Optimizing Storage Assignment in Temperature Zones for Fruits and Vegetables

Mostafa Bahadornia

UHasselt, Research Group Logistics  
e-mail: [mostafa.bahadornia@uhasselt.be](mailto:mostafa.bahadornia@uhasselt.be)

Katrien Ramaekers

UHasselt, Research Group Logistics  
e-mail: [katrien.ramaekers@uhasselt.be](mailto:katrien.ramaekers@uhasselt.be)

Kris Braekers

UHasselt, Research Group Logistics  
e-mail: [kris.braekers@uhasselt.be](mailto:kris.braekers@uhasselt.be)

Trijntje Cornelissens

University of Antwerp, Department of Engineering Management  
e-mail: [trijntje.cornelissens@uantwerpen.be](mailto:trijntje.cornelissens@uantwerpen.be)

Ruben D'Haen

UHasselt, Research Group Logistics  
e-mail: [ruben.dhaen@uhasselt.be](mailto:ruben.dhaen@uhasselt.be)

Food loss in retail and consumer stages accounts for a substantial share of the total waste in fruit and vegetable (F&V) supply chains, while maintaining product freshness requires significant refrigeration energy. Motivated by the rapid expansion of micro-fulfillment centers (MFCs) and the lack of integrated decision-making tools for temperature-controlled storage, this study develops an optimization model for storage location assignment (SLA) of F&V that jointly considers cooling energy consumption, product deterioration in the warehouse, and expected food loss at the consumer level. Unlike existing SLA approaches that focus solely on matching products to ideal storage conditions, our model captures the full freshness–energy trade-off using life cycle assessment as a unified metric for environmental impact. A novel continuous-time representation of the First-Expire-First-Out rule is introduced, enabling explicit estimation of remaining shelf-life and dependencies across batches while avoiding the large set of intermediate variables typical in prior formulations. The proposed MIQP model optimizes the allocation of limited temperature-zone capacity and allows transfers of existing inventory across zones. Then, the proposed model will be validated through an experimental design. The proposed framework provides MFCs with a holistic decision support tool that aligns inventory placement with corporate sustainability goals by reducing both energy use and total food loss across the supply chain.

# Real-time railway timetable rescheduling in case of partial blockages

Inneke Van Hoeck

KU Leuven, Institute for Mobility - CIB

e-mail: [inneke.vanhoeck@kuleuven.be](mailto:inneke.vanhoeck@kuleuven.be)

Pieter Vansteenwegen

KU Leuven, Institute for Mobility - CIB

e-mail: [pieter.vansteenwegen@kuleuven.be](mailto:pieter.vansteenwegen@kuleuven.be)

During railway operations, unexpected incidents such as malfunctioning infrastructure or rolling stock can cause conflicts in the timetable. We specifically consider disruptions that cause a partial blockage of a corridor for several hours. When such a disruption occurs, the dispatchers need to resolve the conflicts by rescheduling the timetable in real-time. In practice, predefined contingency plans are sometimes used to do this. However, this is a very static concept that is not always suitable to handle the dynamic and unique nature of a disruption. More flexible OR-based approaches that are easy to use for the dispatchers are required, as discussed by Cacchiani et al. (2014) and Ghaemi et al. (2017) in their survey papers on this topic.

We propose a novel approach to disruption handling that is based on the work of Van Aken et al. (2019). Despite the unique nature of each disruption, they point out that there does exist some structure in how to handle a disruption. In a railway network, there are certain infrastructure layouts that occur frequently, for example a double-track corridor between double switches. We refer to such a frequently occurring infrastructure layout as an Archetypical Infrastructure Piece, or AIP for short. It has a fixed structure that can be completely described by only a few parameters. A limited number of AIPs can be identified such that they cover a large part of the network. The general idea is then that, even though each disruption is unique in terms of its exact location and duration, a disruption on an AIP can always be handled with a similar approach. For example, if one track of a double-track corridor is blocked, then the solution strategy always comes down to determining in which order the trains can use the remaining track and, possibly, which trains to cancel. Thus, for each AIP, an appropriate optimization model can be developed beforehand to reschedule the timetable. When a disruption occurs, this model only requires some parameter values, that can be determined easily, to be applicable for the specific disruption. This way, many possible disruption scenarios can be handled with a limited number of models. In the optimization models, the AIP is considered in microscopic detail and the possible dispatching measures are retiming, reordering, cancellation and stop-skipping. The objective is to minimize the impact of the disruption on the

passengers.

Due to the inherently uncertain nature of railway operations, it is important that uncertainties are also considered when handling disruptions. Therefore, the optimization models developed for the AIPs are embedded in a rolling horizon approach to dynamically reschedule the timetable. In literature, this approach has already been proposed to handle the uncertain duration of a disruption (e.g., Zhu and Goverde, 2020). Additionally, it allows us to include information about the current status of the timetable during the optimization. More specifically, it is very likely that some trains in the timetable are already driving with a (small) delay before they even reach the disrupted area. By using a rolling horizon approach, current and expected delays can be taken into account when rescheduling the timetable.

We perform experiments for two common AIPs: a double-track corridor between double switches and a more general multi-track corridor between switch areas. Instances with different parameter values for the AIPs, different timetables and multiple delay scenarios are considered. The experiments show that the rolling horizon approach obtains high-quality solutions. Additionally, considering the current train delays during the rescheduling leads to a significant decrease, ranging from 12% to 30%, in the total train delay compared to simply considering the planned timetable.

## References

- Cacchiani, V., Huisman, D., Kidd, M., Kroon, L., Toth, P., Veelenturf, L., Wagenaar, J. (2014). An overview of recovery models and algorithms for real-time railway rescheduling. *Transportation Research Part B: Methodological*, 63:15–37.
- Ghaemi, N., Cats, O., Goverde, R. (2017). Railway disruption management challenges and possible solution directions. *Public Transport*, 9:343–364.
- Van Aken, S., Van Thielen, S., Vansteenwegen, P. (2019). Optimal Timetables in case of Temporarily Unavailable Tracks. In *Online proceedings of 8th International Conference on Railway Operations Modelling and Analysis (ICROMA), RailNorrköping*.
- Zhu, Y., Goverde, R. (2020). Dynamic and robust timetable rescheduling for uncertain railway disruptions. *Journal of Rail Transport Planning & Management*, 15:100196.

# Do profit-driven learning methods deliver profit? A policy evaluation study

Shimanto Rahman

Ghent University, Faculty of Economics and Business Administration  
e-mail: [shimanto.rahman@ugent.be](mailto:shimanto.rahman@ugent.be)

Bram Janssens

Ghent University, Faculty of Economics and Business Administration  
e-mail: [bram.janssens@ugent.be](mailto:bram.janssens@ugent.be)

Matthias Bogaert

Ghent University, Faculty of Economics and Business Administration  
e-mail: [matthias.bogaert@ugent.be](mailto:matthias.bogaert@ugent.be)

Machine learning is increasingly applied in domains where decision outcomes carry heterogeneous economic consequences. In credit scoring, fraud detection, and customer churn prediction [1], the goal is not merely predictive accuracy but financial impact. This has motivated research on cost-sensitive and profit-driven learning, each proposing tailored metrics, loss functions, and models intended to maximize business value. Yet despite this methodological richness, empirical studies overwhelmingly evaluate these approaches on static datasets, implicitly assuming that optimizing cost- or profit-based metrics leads to superior real-world decision policies. Whether this assumption holds remains largely untested.

This research addresses this gap by shifting the focus from predictive evaluation to policy evaluation. Rather than asking which method scores best on Expected Cost or Expected Maximum Profit, the aim is to examine whether these methods yield stable, reliable, and value-generating decisions when deployed over time. Because ground-truth probabilities, true costs, and data-generating mechanisms are unknown in real-world settings, the study employs controlled synthetic environments. Synthetic data enables explicit cost structures, tunable instance-dependent mechanisms, and designed distribution shifts, allowing conjectures about stability, robustness, and optimality to be explored systematically. As of writing, empirical results are still in progress, but the framework is fully established.

A central line of inquiry concerns the role of properness in value-based metrics. Recent work by Xu & Ergu [2] shows that Expected Cost is not a proper scoring rule: minimizing it does not necessarily yield true conditional probabilities. This raises an open question about whether proper scoring rules (e.g., log-loss) might produce more reliable long-term business outcomes—even if they are not explicitly cost-aligned—or whether improper but economically meaningful metrics remain preferable. This research investigates these competing intuitions through environments where correctness of probability estimation and correctness of de-

cisions can be disentangled.

A second focus is the robustness of instance-dependent cost structures. While instance-dependent costs allow finer modeling of transaction amounts, customer value, or individualized risk [3], recent studies suggest that this granularity may introduce unstable explanations [4] and more volatile learning dynamics [5]. Using synthetic shifts—e.g., changes in customer value distributions or altered cost-generating processes—the research examines the conjecture that instance-dependent costs may lose their advantage when these assumptions drift, potentially making simpler class-dependent cost schemes more reliable for long-term policy decisions.

Finally, this research examines the trade-off between metaheuristic and gradient-based optimization in value-based learning. Both approaches explicitly incorporate costs, but differ in how economic objectives are optimized. Metaheuristic methods directly optimize non-differentiable profit metrics, such as Expected Maximum Profit, at the expense of higher computational cost and extensive hyperparameter tuning. Gradient-based cost-sensitive methods instead optimize differentiable objectives like Expected Cost, and are typically faster, more stable, and require fewer design choices. This research investigates whether the additional complexity of metaheuristic optimization leads to superior policy-level outcomes, or whether gradient-based methods offer a more effective balance between performance, robustness, and practicality.

Taken together, this research seeks to articulate the practical conditions under which value-based learning methods deliver on their promise. By evaluating the policy-level impact of these techniques rather than solely their metric-level performance, this study aims to provide a clearer understanding of when cost-sensitive and profit-driven methods generate robust economic gains, when they falter, and how future methodological advances might better align predictive modeling with real-world decision-making.

## References

- [1] Vanderschueren, T., Verdonck, T., Baesens, B., & Verbeke, W. (2022). Predict-then-optimize or predict-and-optimize? An empirical evaluation of cost-sensitive learning strategies. *Information sciences*, 594, 400-415.
- [2] Xu, Y., & Ergu, D. (2025). Understanding the profit-driven credit scoring: New insights on loss functions and measures. *Applied Soft Computing*, 114184.
- [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*, 341(1), 267-293.
- [4] Ballegeer, M., Bogaert, M., & Benoit, D. F. (2025). Evaluating the stability of model explanations in instance-dependent cost-sensitive credit scoring. *European Journal of Operational Research*.
- [5] De Vos, S., Vanderschueren, T., Verdonck, T., & Verbeke, W. (2023). Robust instance-dependent cost-sensitive classification. *Advances in data analysis and classification*, 17(4), 1057-1079.

# A Lightweight Graph Representation for Scalable Deep Learning on 3D CAD Data

Toon Van Camp

Ghent University, Research Group Data Analytics

FlandersMake@UGent-corelab CVAMO

e-mail: `toon.vancamp@ugent.be`

Dries F. Benoit

Ghent University, Research Group Data Analytics

FlandersMake@UGent-corelab CVAMO

e-mail: `dries.benoit@ugent.be`

Design-for-manufacturing aims to ensure that a product can be produced reliably, efficiently, and at acceptable cost. In practice, manufacturability issues are often discovered only after engineers complete their computer-aided design (CAD) models, leading to costly redesign cycles. This motivates the use of data-driven models that provide early feedback directly from CAD data. Deep learning offers a way to provide earlier, automated feedback by learning from large collections of existing CAD data. These models can support several important tasks. One example is manufacturability analysis, where a model evaluates a design to identify potential production problems, such as overly thin walls or difficult-to-reach areas, that could complicate fabrication or increase costs. Another is part retrieval, which helps engineers find existing designs that are similar to a new one, enabling part reuse and reducing redundant work. However, a central challenge remains how to represent 3D CAD geometry in a form that neural networks can process efficiently.

Early 3D deep learning approaches relied on conversion-based representations such as point clouds [1] and voxel grids [2]. Point clouds contain large sets of sampled surface points without information about how surfaces connect, while voxel grids discretize 3D space into small cubes. These methods often suffer from information loss or high computational overhead during model training. Graph-based representations have emerged as a powerful alternative, allowing for direct operation on the boundary representation (B-rep) of 3D solids. The B-rep is a widely used format in CAD systems, which defines a solid through its geometric elements such as surfaces, edges, and vertices, and describes how these elements are connected. A fundamental deep learning approach for directly processing B-rep data is UV-Net [3]. UV-Net represents a 3D solid as a graph whose nodes correspond to surfaces and whose node connections represent the shared edges between those surfaces. Moreover, it constructs high-dimensional features by sampling points on surfaces and edges. Although effective, these dense samples produce large graphs and computationally-demanding training pipelines.

We introduce EdgeNet, a lightweight graph representation for deep learning

on CAD data. Instead of modeling surfaces, EdgeNet focuses exclusively on the edges and vertices of the B-rep, which can be visualized as the “skeleton” of the design. Each edge in the solid is treated as a graph node, and two nodes are connected when their corresponding edges meet at a vertex. Node features are constructed by sampling points along each edge and recording their 3D coordinates and tangent vectors, providing a compact description of local geometry. Based on the UV-Net architecture, these sampled sequences are then processed by a 1D convolutional network to produce a fixed-length embedding for every node. With these embeddings in place, a Graph Isomorphism Network (GIN) performs message passing to exchange information between adjacent nodes and capture higher-level geometric structure. Hierarchical pooling is then applied to aggregate information across graph layers, ultimately producing a single 1D graph-level embedding that summarizes the entire CAD model. This embedding can then be passed to a fully connected neural network or any subsequent prediction module, depending on the downstream task. For self-supervised learning, EdgeNet uses a contrastive learning framework with graph-specific augmentations such as node removal and feature masking to create positive samples, enabling scalable pretraining without labeled data.

Across four classification benchmarks, EdgeNet consistently matches or surpasses the performance of heavier graph-based models while substantially reducing computational requirements. It achieves up to  $3.3\times$  smaller model size, accelerates training by up to  $2.65\times$ , and reduces storage requirements by up to  $9\times$ . In a self-supervised setting, EdgeNet also demonstrates strong retrieval performance in similarity search tasks, highlighting its potential for scalable pretraining and CAD model retrieval. These results indicate that an edge-only representation provides sufficient information for both classification and retrieval tasks, while delivering significant gains in efficiency and scalability.

## References

- [1] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [2] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 922–928. Ieee, 2015.
- [3] Pradeep Kumar Jayaraman, Aditya Sanghi, Joseph G Lambourne, Karl DD Willis, Thomas Davies, Hooman Shayani, and Nigel Morris. Uv-net: Learning from boundary representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11703–11712, 2021.

# FoV-Net: Rotation-Invariant 3D Deep Learning for CAD Analysis

Matteo Ballegeer<sup>1</sup>, Dries F. Benoit<sup>1</sup>,

<sup>1</sup>Ghent University, Research group Data Analytics,

<sup>2</sup>FlandersMake@UGent – core lab CVAMO, Belgium

`matteo.ballegeer@UGent.be`, `dries.benoit@UGent.be`

Automated analysis of 3D designs is becoming a cornerstone of modern manufacturing, supporting applications such as part classification, manufacturability assessment, and quality control. Industrial designs are typically stored as Computer-Aided Design (CAD) models, which use boundary representations (B-reps) to describe parts through their faces, edges, and vertices. Recent progress in deep learning has shown that directly learning from these B-reps can greatly improve the fidelity of 3D model understanding. However, existing methods suffer from a critical limitation: they rely heavily on global coordinates and orientations. As a result, even a simple rotation of a CAD model, common in real-world pipelines where files arrive in arbitrary orientations, can cause the learned features to drift and severely degrade performance.

We introduce FoV-Net, a new geometric descriptor and deep learning framework designed to overcome this limitation and deliver rotation-invariant analysis of 3D CAD data. Instead of depending on global geometry, FoV-Net represents each face using two complementary, inherently rotation-robust descriptors. The first, a Local Reference Frame (LRF) UV-grid, captures the detailed shape of each face within a coordinate system defined locally and consistently across rotations. The second, a Field-of-View (FoV) descriptor, applies ray casting to summarize the surrounding 3D context from the perspective of each face, capturing structural context. Because both descriptors are anchored locally, their representations remain stable regardless of the part's orientation.

FoV-Net extracts per-face features from these descriptors using lightweight CNNs and integrates them across the full part through a graph neural network that mirrors the connectivity of the B-rep. Evaluated on several benchmark datasets, Our approach achieves state-of-the-art performance on both part classification and face segmentation, especially under arbitrary rotations. Moreover, the method requires substantially less training data than previous methods—a crucial advantage in industrial settings where labeled data is often scarce. By directly addressing orientation sensitivity and aligning with the realities of industrial CAD workflows, FoV-Net takes a significant step toward reliable, deployable deep learning for 3D CAD understanding.

# A coevolutionary approach for optimizing interpretable rule ensembles

Simon De Lange

Ghent University, Department of Marketing, Innovation and Organisation  
FlandersMake@UGent–corelab CVAMO, Belgium  
e-mail: [simondla.delange@ugent.be](mailto:simondla.delange@ugent.be)

Matthias Bogaert

Ghent University, Department of Marketing, Innovation and Organisation  
FlandersMake@UGent–corelab CVAMO, Belgium  
e-mail: [matthias.bogaert@ugent.be](mailto:matthias.bogaert@ugent.be)

Koen W. De Bock

Audencia Business School, Department of Marketing  
e-mail: [kdebock@audencia.com](mailto:kdebock@audencia.com)

Dirk Van den Poel

Ghent University, Department of Marketing, Innovation and Organisation  
FlandersMake@UGent–corelab CVAMO, Belgium  
e-mail: [dirkvandenpoel@ugent.be](mailto:dirkvandenpoel@ugent.be)

Model interpretability has become an increasingly important aspect of machine learning, particularly in high-stakes domains such as healthcare, finance, and criminal justice, where understanding the decision-making process is crucial for trust, accountability, and regulatory compliance (Rudin, 2019). To balance predictive performance with interpretability, various methods have been proposed, including rule-based models (e.g., decision trees) and linear models. More recently, rule ensembles have gained popularity as an interpretable machine learning method that linearly combines an ensemble of rules with the original features to obtain an interpretable model. In business analytics, several variants have been shown to achieve competitive predictive performance while maintaining interpretability (De Bock & De Caigny, 2021). The key challenge in constructing rule ensembles lies in identifying a compact and relevant set of rules from a potentially vast search space that, when linearly combined, yields strong predictive performance. The predominant approach to building rule ensembles is to first fit a tree ensemble to determine a large set of rule candidates and subsequently fit a regularized linear model to obtain a sparse, additive, and interpretable model that achieves competitive predictive performance by allowing for non-linearities and interactions between features (Friedman & Popescu, 2008).

However, this two-step approach is greedy in nature and does not guarantee an optimal model. First, tree ensembles are typically designed to maximize predictive performance by combining multiple weak learners, rather than to extract

a compact and relevant set of rules for a linear model. Second, the decision trees are trained without considering the final linear model. Therefore, the tree ensemble will attempt to approximate linear relationships using multiple rules, which could be more efficiently captured by a single linear term.

To address these limitations of traditional rule ensemble methods, we propose CEGARE, a CoEvolutionary Genetic Algorithm Rule Ensemble that simultaneously optimizes the rules and the coefficient vector of the rule ensemble. By employing a coevolutionary genetic algorithm (Potter & De Jong, 1994), CEGARE evolves a population consisting of multiple species of candidate rules alongside the coefficients of the linear model, allowing for a more holistic optimization process. To validate CEGARE, we conduct a benchmark on multiple datasets from various business domains to compare predictive performance and sparsity against conventional rule ensembles. Preliminary results indicate that CEGARE achieves competitive accuracy while producing significantly more compact and interpretable models, demonstrating its potential as a viable alternative to existing rule ensemble approaches.

## References

- De Bock, K. W., & De Caigny, A. (2021). Spline-rule ensemble classifiers with structured sparsity regularization for interpretable customer churn modeling. *Decision Support Systems*, 150, 113523. <https://doi.org/https://doi.org/10.1016/j.dss.2021.113523>
- Friedman, J. H., & Popescu, B. E. (2008). Predictive learning via rule ensembles. *The Annals of Applied Statistics*, 2(3), 916–954. <https://doi.org/10.1214/07-AOAS148>
- Potter, M. A., & De Jong, K. A. (1994). A cooperative coevolutionary approach to function optimization. In Y. Davidor, H.-P. Schwefel, & R. Männer (Eds.), *Parallel Problem Solving from Nature — PPSN III* (pp. 249–257). Springer Berlin Heidelberg.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.

# Trusted Triplet Contrastive Learning for Multi-view clustering

Yan Li

Ghent University,

Faculty of Economics and Business Administration,

Research Group Data Analytics

FlandersMake@UGent, corelab CVAMO

e-mail: `Yan.Li@UGent.be`

Dries F. Benoit

Ghent University,

Faculty of Economics and Business Administration

FlandersMake@UGent, core lab CVAMO

e-mail: `Dries.Benoit@UGent.be`

Multi-view clustering aims to exploit consistent and complementary information from multiple views to uncover underlying cluster structures [1]. Recently, contrastive learning [2] has emerged as a powerful tool for unsupervised representation learning and has been introduced into multi-view clustering, giving rise to multi-view contrastive learning (MVCL). By aligning the representations of the same instance across views, MVCLs have achieved promising improvements. Most approaches increasingly adopt contrastive learning to enforce cross-view consistency at the instance [3] or cluster level [4]. However, existing MVCL methods still suffer from two fundamental limitations.

First, their alignment objectives are largely semantic-agnostic: consistency is enforced based on instance pairing, without modeling semantic relationships among different instances. As a result, intra-view semantic cohesion is weak, semantic false negatives are introduced in cross-view contrast [5], and unreliable cluster assignments further distort cluster-level alignment across views. Second, most methods implicitly assume that all views are equally reliable [6], overlooking view-wise uncertainty and quality heterogeneity. Consequently, noisy or weak views are treated on par with informative ones, allowing unreliable information to propagate through cross-view alignment and degrade final clustering performance.

To address these issues, we propose a novel Trusted Triplet Contrastive Learning (TTCL) for Multi-view clustering. By explicitly modeling semantic evidence and its associated uncertainty, our method produces more reliable pseudo-labels to guide triplet contrastive learning at multiple levels. Specifically, semantic evidence is used to construct more reliable self-pseudo-label graphs, which enable semantic-aware intra-view weighted contrastive learning. Pseudo-label graphs further help identify semantically similar samples and filter out false negatives

in cross-view instance-level contrast. In addition, evidence-based representations are leveraged for cross-view cluster-level contrast to improve the cluster consistency. Beyond triplet alignment, we further model view-wise uncertainty and employ uncertainty-aware fusion to reweight view-specific clustering results, thereby reducing the influence of unreliable views in the final decision. This unified framework allows semantic evidence and uncertainty to jointly enhance alignment reliability and multi-view fusion.

Extensive experiments are conducted on multiple widely used MVCL benchmarks to comprehensively evaluate the proposed framework. The results demonstrate consistent and significant performance improvements over state-of-the-art methods across different data characteristics and view configurations. In addition, thorough ablation studies are performed to disentangle the contributions of each component, verifying the effectiveness of semantic evidential learning, triplet contrastive alignment, and uncertainty-aware fusion. These analyses confirm that the proposed triplet contrastive learning strategy and evidence-based uncertainty modeling jointly play a crucial role in improving alignment reliability and multi-view clustering robustness.

## References

- [1] J. Chen, H. Mao, W. L. Woo, X. Peng, Deep multiview clustering by contrasting cluster assignments, *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 16752–16761.
- [2] K. He, H. Fan, Y. Wu, S. Xie, R. Girshick, Momentum contrast for unsupervised visual representation learning, *arXiv:1911.05722*, 2020.
- [3] Y. Tian, D. Krishnan, P. Isola, Contrastive multiview coding, *European conference on computer vision (ECCV)*, 2020, pp. 776–794.
- [4] J. Xu, H. Tang, Y. Ren, L. Peng, X. Zhu, L. He, Multi-level feature learning for contrastive multi-view clustering, *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 16030–16039.
- [5] S. Hu, B. Tian, W. Liu, Y. Ye, Self-supervised trusted contrastive multi-view clustering with uncertainty refined, *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025, pp. 17305–17313.
- [6] Z. Han, C. Zhang, H. Fu, J. T. Zhou, Trusted multi-view classification with dynamic evidential fusion, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45 (2) (2022) 2551–2566.

# Personalized Advertising with Generative AI through Interpretable User–Creative Clusters

Nathan Marckx<sup>1</sup>, Dylan Van Mulders<sup>1</sup>,  
Matthias Bogaert<sup>1</sup>, Dirk Van den Poel<sup>1</sup>

<sup>1</sup>Ghent University, Department of Marketing, Innovation and Organisation,  
FlandersMake@UGent – core lab CVAMO, Belgium

Nathan.Marckx@UGent.be, Dylan.VanMulders@UGent.be,  
Matthias.Bogaert@UGent.be, Dirk.VandenPoel@UGent.be

**Keywords:** digital advertising, multimodal modeling, creative optimization

Digital advertising has traditionally optimized *who* to target through audience segmentation and behavioral recommendation systems [1]. These systems typically rely on third-party user-level data and evaluate performance through impressions, clicks, and click-through rate (CTR), defined as the share of users who click an ad after viewing it. However, they often treat ad creatives (image, message, and layout) as static identifiers, overlooking the actual visual and textual content [2, 3]. As a result, marketers gain limited insight into *what* content should be shown or why certain styles succeed with specific audiences, reinforcing black-box behavioral targeting practices. With generative AI (GenAI) increasingly shaping creative workflows, this lack of interpretability around creative influence becomes a critical limitation for effective personalization.

The emergence of generative models capable of dynamic content creation and visual adaptation creates new opportunities [4, 5], but effective use requires a transparent connection between audience characteristics and creative styles. Without such grounding, generative systems lack the necessary signals to produce audience-relevant outputs. For example, a model prompted to “appeal to young families” may generate irrelevant or generic imagery, or hallucinate attributes not present in the data. In addition, while audience clustering and CTR-based creative grouping exist, few approaches jointly cluster audiences and creatives in an interpretable manner. This gap motivates a structured methodology in which transparent, data-driven clusters on both sides are aligned through observed user–creative interaction patterns to reveal which creative attributes resonate with which audience segments. Under such a structure, GenAI can be guided to generate creatives personalized at the audience-segment level in a controllable and explainable manner, moving beyond black-box behavioral targeting.

To address this need, an interpretable dual-clustering framework is proposed using creative impression–click data from a large Belgian advertising company. Audience segments are formed using available database features, including first-party sociodemographics, device type, and third-party interests. Creative clusters, in contrast, are constructed from engineered multimodal features including

headline sentiment (via RobBERT v2), low-level visual descriptors (color, contrast, edge density) [6], object detections (YOLOv11s), and high-level scene, emotion, and style descriptors generated by the multimodal Gemma 3 model. Additionally, joint CLIP embeddings of headline–image pairs are projected via UMAP to form semantic representations. Both user and creative clusters are selected through an automated model selection pipeline comparing K-Means, Gaussian Mixture Models, hierarchical clustering, and HDBSCAN, evaluated using standard internal validity indices and custom complexity penalties.

The selected models—K-Means for audiences (21 clusters) and hierarchical clustering for creatives (20 clusters)—support the construction of a CTR-based user–creative interaction matrix that exposes audience–creative alignment. Results show, for example, that creatives with positive headlines, warm color palettes, and benefits-oriented visuals overperform among older, higher-income smartphone users. Another creative cluster, featuring outdoor scenes, gardening visuals, neutral moods, and animals, achieved high CTR across nearly all segments yet remained underexposed, indicating substantial untapped potential.

To demonstrate operational value, the framework is extended with a GenAI personalization layer. For each user cluster, the system identifies top-performing creative clusters and converts their dominant features into structured prompts using Gemma 3, which are then supplied to the Flux Schnell diffusion model to generate visually aligned creatives. This application illustrates how interpretable clustering can steer controllable and transparent generative personalization.

The proposed framework provides a clear, actionable pathway from audience segmentation to creative development. By aligning interpretable user and creative clusters, it supports scalable, human-supervised generative personalization and offers a transparent alternative to black-box behavioral targeting in modern digital advertising.

## References

- [1] B. Gao, Y. Wang, H. Xie, Y. Hu, and Y. Hu, *Artificial Intelligence in Advertising: Advancements, Challenges, and Ethical Considerations in Targeting, Personalization, Content Creation, and Ad Optimization*, SAGE Open, vol. 13, no. 4, 2023.
- [2] W. Zhang, P. Zhang, J. Dong, Y. Wang, P. Zhang, B. Zhang, X. Wang, and D. Wang, *Cross-Element Combinatorial Selection for Multi-Element Creative in Display Advertising*, arXiv:2307.01593, 2023.
- [3] J. Chen, T. Ge, G. Jiang, Z. Zhang, D. Lian, and K. Zheng, *Efficient Optimal Selection for Composited Advertising Creatives with Tree Structure*, arXiv:2103.01453, 2021.
- [4] X. Lu, *Adaptive Optimization of Advertising Creative Visual Elements Based on Multi-dimensional User Behavior Data*, in Proceedings of the International Conference on Generative Artificial Intelligence for Business, 2025, pp. 360–368.
- [5] Q. Gu, Y. Li, D. Zeng, L. Wang, M. Pang, C. Peng, Z. Lin, C. Law, and J. Shao, *Generative Modeling with Multi-Instance Reward Learning for E-commerce Creative Optimization*, arXiv:2508.09730, 2025.
- [6] J. Azimi, R. Zhang, Y. Zhou, V. Navalpakkam, J. Mao, and X. Fern, *The impact of visual appearance on user response in online display advertising*, in Proceedings of the 21st International Conference on World Wide Web (WWW Companion), 2012, pp. 457–458.

# An Empirical Exploration of Explanation-Guided Learning for Causal Effect Estimation

Anna Glado

University of Padova  
e-mail: [anna.glado@vub.be](mailto:anna.glado@vub.be)

Luc Hirsch

Vrije Universiteit Brussels, Data Lab  
e-mail: [luc.hirsch@vub.be](mailto:luc.hirsch@vub.be)

Alessandro Marchese

Vrije Universiteit Brussels, Data Lab  
e-mail: [alessandro.marchese@vub.be](mailto:alessandro.marchese@vub.be)

Sam Verboven

Vrije Universiteit Brussels, Data Lab  
e-mail: [sam.verboven@vub.be](mailto:sam.verboven@vub.be)

## 1 Intro

Estimating causal effects from observational data remains challenging due to limited data availability and the presence of complex confounding relationships between variables. Yet, existing causal inference methods typically ignore human-generated explanations that encode additional information about the mechanisms underlying decisions and outcomes. In this paper, we explore how externally provided, human-generated explanations can improve causal effect estimation. These explanations are articulated by humans—such as domain experts, decision-makers, or users—and arise naturally in many operational systems. Examples include reasons recorded in resource allocation mechanisms, reported motives for product returns, stated reasons for customer churn, or structured user feedback in recommender systems. Building on results in organ allocation studies [3], which demonstrate that explanations can meaningfully improve treatment effect estimation, we aim to provide broader empirical insights into explanation-guided causal machine learning.

## 2 Problem Setting

We investigate Explanation-Guided Learning (EGL) [1, 6, 5] as a model-agnostic complement to existing causal methods. Here, ‘explanations’ refer to expert-provided rationales or feature-level annotations, and EGL denotes training procedures that incorporate such explanation signals—together with model-generated explanations—as auxiliary supervision to steer the causal estimator. This leads to the following formulation of a minimisation problem:

$$\min_f \mathcal{L}(f(X, T), Y) + \mathcal{L}(g(f, X, T, Y), E).$$

Here,  $f$  denotes the causal estimator,  $X, T, Y$  the covariates, treatment, and outcome, and  $E$  the explanation representation. The explainer  $g$  is a model

that maps  $(f, X, T, Y)$  to the model-generated explanations. A specific type of explainer is chosen based on the format of human-given explanations to ensure their alignment is feasible.

### 3 Contribution

Moving beyond direction-only explanations used in [3], we explore more varied explanation types that help the model learn the underlying causal mechanism. These explanations encode which features most strongly influenced the treatment or outcome assignment in each instance. To the best of our knowledge, we are the first to examine how such explanations can be leveraged within a causal estimation framework. Since existing explainers were developed primarily for predictive interpretability rather than causal inference, it is unclear a priori which of them provide signals that are useful for estimating causal effects. We synthetically evaluate three popular explainers with matching ground-truth explanations based on current EGL practice [1]: LIME [4], which approximates the model locally using a simple surrogate; counterfactual explanations [1], which identify minimal input changes that alter a prediction; and SHAP [2], which attributes feature contributions via Shapley-value principles. Understanding how different expert and model explanation combinations behave under controlled conditions sheds light on when explanation-guided methods can meaningfully improve causal effect estimation.

### References

- [1] Yuyang Gao, Siyi Gu, Junji Jiang, Sungsoo Ray Hong, Dazhou Yu, and Liang Zhao. Going beyond xai: A systematic survey for explanation-guided learning. *ACM Computing Surveys*, 56(7):1–39, 2024.
- [2] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [3] Alessandro Marchese, Jeroen Berrevoets, and Sam Verboven. Causal explanation-guided learning for organ allocation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [4] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.
- [5] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- [6] Leander Weber, Sebastian Lopuschkin, Alexander Binder, and Wojciech Samek. Beyond explaining: Opportunities and challenges of xai-based model improvement. *Information Fusion*, 92:154–176, 2023.

# POSEIDON: Planned Operations and Scheduling Engine for Integrated Demining in Oceanic Navigation

Philippe Blondeel  
 Royal Military Academy, Belgium  
 Department of mathematics  
 e-mail: philippe.blondeel@mil.be

Filip Van Utterbeeck  
 Royal Military Academy, Belgium  
 Department of mathematics  
 e-mail: filip.vanutterbeeck@mil.be

Ben Lauwens  
 Royal Military Academy, Belgium  
 Department of mathematics  
 e-mail: ben.lauwens@mil.be

Planning mine countermeasure (MCM) missions at sea is a challenging and time-sensitive task that traditionally relies on human expertise. To automate and optimize this process, we present our proof of concept planning and scheduling engine: POSEIDON. Our engine uses advanced mathematical optimization techniques, such as stochastic optimal control and combinatorial optimization, in order to integrally plan and schedule a mine countermeasure mission. The engine consists of two interconnected components: a low-level trajectory planner and a high-level mission scheduler. The output of the engine consists of an optimized schedule for planning mine countermeasure missions at sea.

## Low-Level component

The low-level component is responsible for computing paths for autonomous underwater vehicles (AUVs) operating in user-defined zones, such that the residual risk of mine presence after surveying is below a user-defined threshold. Two path planning strategies are possible in the engine: a classical boustrophedon search pattern and a stochastic optimal control approach see [1, 2]. Our engine allows for the computation of trajectories of multiple AUVs operating concurrently in a generalized convex polygonal domain, e.g., quadrilateral, pentagon, hexagon, heptagon. An output example of a trajectory for a quadrilateral domain is shown in Fig. 1.

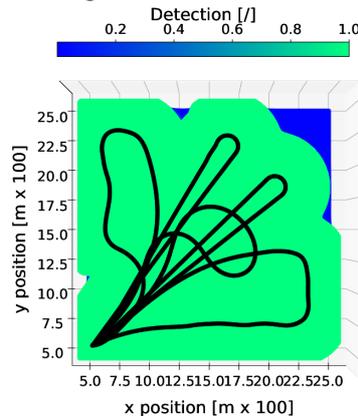


Figure 1: Trajectories of five AUVs in a quadrilateral domain where green stands for the surveyed area and blue for the unsurveyed area. The remaining residual risk is below 5%.

## High-Level component

The goal of the high-level component is then to construct a mission schedule that minimizes total operation time based on the simulated times originating from the low-level planner. The combinatorial scheduling problem of the different actions of the actors is formulated as a Mixed Integer Linear Programming (MILP) model, and was first proposed in [3]. A schedule example is shown in Fig. 2 where we consider one zone, two service agents and one transport agent.

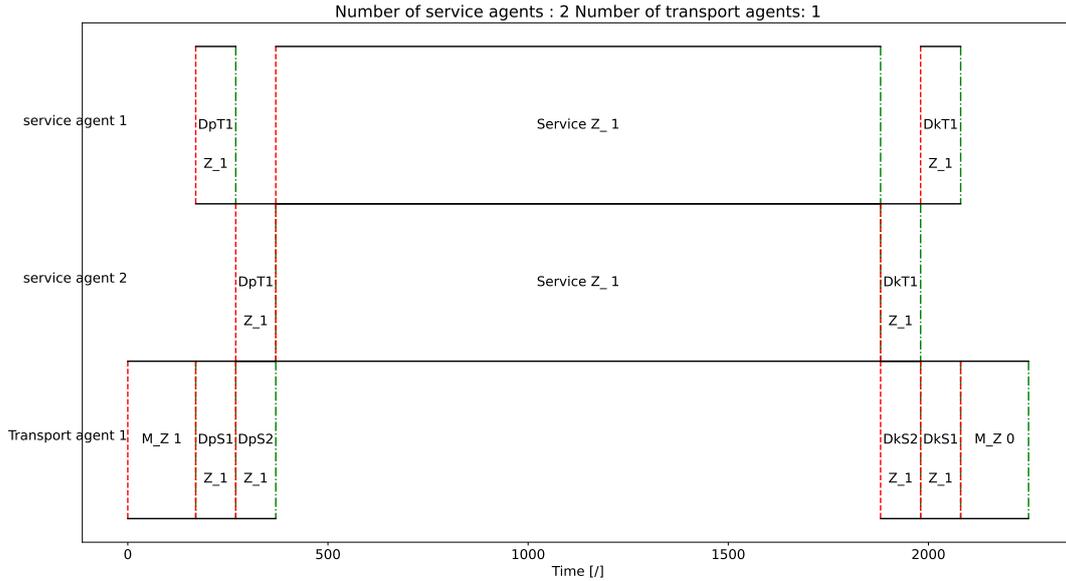


Figure 2: Example of a mine countermeasure schedule operation with one zone, two service agents and one transport agent. DpT1 stands for a deploy action executed by transport agent 1 while DkT1 is a dock action executed by transport agent 1, both applied to a specific service agent. DpS1 stands for a deploy action applied to service agent 1 while DkS1 is a dock action applied to service agent 1, both executed by a specific transport agent. M\_Z\_1 stands for a move action to zone 1 executed by any applicable agent. Service Z\_1 is a survey action executed by any applicable service agent.

## References

- [1] P. Blondeel, F. Van Utterbeeck, and B. Lauwens, “Modeling sand ripples in mine countermeasure simulations by means of stochastic optimal control”, in *The 19th European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS, Lisbon, Portugal, 2024*.
- [2] P. Blondeel, F. V. Utterbeeck, and B. Lauwens, “Application of quasi-monte carlo in mine countermeasure simulations with a stochastic optimal control framework”, *Proceedings of the 6th International Conference on Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing (To be published), 2026*.
- [3] M. J. Bays and T. A. Wettergren, “Service agent–transport agent task planning incorporating robust scheduling techniques,” *Robotics and Autonomous Systems*, vol. 89, pp. 15–26, 2017.

# Scheduling in-orbit servicing missions for mega-constellations

Sushanta Nigudkar

University of Antwerp, Department of Engineering Management

e-mail: [Sushanta.Nigudkar@uantwerpen.be](mailto:Sushanta.Nigudkar@uantwerpen.be)

Christof Defryn

University of Antwerp, Department of Engineering Management

e-mail: [christof.defryn@uantwerpen.be](mailto:christof.defryn@uantwerpen.be)

Over the past two decades, modern communication and space data companies have ushered-in the era of mega-constellations: arrangements of thousands of satellites that move in unison and function as a single network (e.g., Starlink from SpaceX or Project Kuiper from Amazon). The total number of active satellites in low-earth orbit (160-2000 kms above the Earth's surface) has grown from a few hundred to over 10,000 satellites over the last 20 years [2]. As this number is expected to grow more rapidly over the coming years, there is a need for active steps to manage in-space assets and guarantee their service. In-orbit servicing is put forward as a promising area to realise this.

In-orbit services can consist of monitoring missions that assess satellite health, refuelling and orbit raising missions that extend satellite lifetimes, and impulsive collision avoidance missions that mitigate the creation of space debris and allow spacecraft to operate safely over their full service lifetime. Research contributions that allow the planning of fuel and time optimal missions to service mega-constellations in low-earth-orbit could have a large impact on ensuring the sustainable use of space over the coming years [3,1].

My work explores this domain through a candidate scenario where a permanent in-orbit service depot is established, from which service shuttles can be dispatched to client satellites that are part of a mega-constellation. In this construct, two problems are investigated: (i) the scheduling of service missions to allow for the time optimal execution of servicing tasks, and (ii) the selection of fuel optimal trajectories between the depot and client satellites.

The complex space environment coupled with orbital motions of both client and servicing satellites makes these problems more challenging than their on-Earth equivalents. Fuel costs cannot be represented by closed-form analytical expressions and are not linearly dependent on time, the design of orbital trajectories cannot take advantage of standard pre-defined paths (e.g. roads on the Earth) and physical phenomenon such as the gravitational effects of the bulge of the earth can have a large impact on fuel efficiency and must therefore be carefully included.

We tackle these complexities using a custom-built simulation environment to test and build optimisation algorithms. A semi-analytical method (Lambert

transfers) is used to estimate fuel costs, and a local search based heuristic is developed to construct fuel-optimal schedules. Moreover, a dynamic program that leverages the dependence of fuel costs on the shapes of transfer orbits and times of flight is also developed.

Future work is envisioned to integrate more exact, indirect trajectory optimisation methods — such as Q-laws — to construct optimal control laws that can ensure collision avoidance directly via the thrust direction at each time step.

## References

- (1) Adam Abdin. “Strategic Management of On-Orbit Servicing: Leveraging Operations Research Methods for Enhanced Mission Planning and Scheduling”. *18th International Conference on Space Operations*. (2025).
- (2) ESA Space Debris Office. “ESA Space Environment Report 2025”. [https://www.esa.int/Space\\_Safety/Space\\_Debris/ESA\\_Space\\_Environment\\_Report\\_2025](https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025). (Accessed: 2025-11-25).
- (3) Tristan Sarton du Jonchay et al. “Framework for modeling and optimization of on-orbit servicing operations under demand uncertainties”. *Journal of Spacecraft and Rockets* 58.4 (2021) pp. 1157–1173.

## Scheduling a single parallel-batching machine with Run-to-Run constraints

Fan Yang

Shanghai Normal University, School of Finance and Business

e-mail: [fan\\_yang@shnu.edu.cn](mailto:fan_yang@shnu.edu.cn)

Xiajie Yi

IESEG School of Management

e-mail: [x.yi@ieseg.fr](mailto:x.yi@ieseg.fr)

Hongtao Wang

University of Southampton, School of Mathematical Sciences

e-mail: [hw5n25@soton.ac.uk](mailto:hw5n25@soton.ac.uk)

Dries Goossens

Ghent University, Faculty of Economics and Business Administration

e-mail: [dries.goossens@ugent.be](mailto:dries.goossens@ugent.be)

Roel Leus

KU Leuven, Faculty of Economics and Business

e-mail: [roel.leus@kuleuven.be](mailto:roel.leus@kuleuven.be)

This study addresses a scheduling problem arising in semiconductor manufacturing, where run-to-run (R2R) control constraints must be considered. We examine a single parallel-batching machine that processes incompatible job families with non-identical job sizes. R2R controllers are typically employed in these settings to adjust control actions or recipes. The machine can simultaneously process multiple jobs from the same family, provided the machine capacity is not exceeded. All jobs within a batch begin and finish processing simultaneously. The objective is to minimize the total weighted completion time. Using the three-field notation introduced by [1], we denote this problem as  $1|p - batch, v_i, incompatible, R2R|\sum w_i C_i$ , where  $p - batch$  indicates a parallel-batching machine,  $v_i$  represents the size of job  $i$ ,  $incompatible$  signifies incompatible job families,  $R2R$  denotes the R2R control constraints, and  $\sum w_i C_i$  represents the total weighted completion time. Figure 1 illustrates a typical R2R control framework, where the R2R controller adjusts the recipe for the next run based on the measurement data collected from the metrology system after processing a blank job/batch.

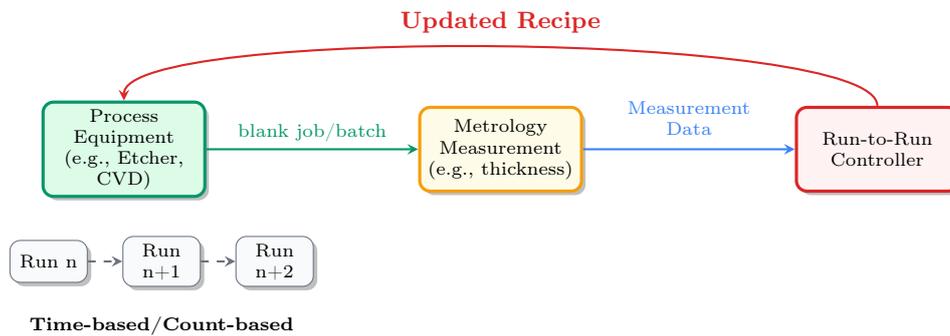


Figure 1: A Run-to-Run Control Framework

Some literature about machine scheduling and R2R constraints can be found in [2], [3], [4], [5], and [6]. To solve this problem, two mathematical programming models, a column generation-based heuristic, and an iterated greedy method are proposed in this study.

## References

- [1] Graham, R. L., Lawler, E. L., Lenstra, J. K., & Rinnooy Kan, A. H. G. (1979). Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of Discrete Mathematics*, 5, 287–326.
- [2] Li, L., Qiao, F., & Wu, Q. (2009). Aco-based multi-objective scheduling of parallel batch processing machines with advanced process control constraints. *The International Journal of Advanced Manufacturing Technology*, 44, 985–994.
- [3] Cai, Y., Kutanoglu, E., Hasenbein, J., & Qin, J. (2012). Single-machine scheduling with advanced process control constraints. *Journal of Scheduling*, 15, 165–179.
- [4] Obeid, A., Dauzère-Pérès, S., & Yugma, C. (2014). Scheduling job families on non-identical parallel machines with time constraints. *Annals of Operations Research*, 213(1), 221–234.
- [5] Nattaf, M., Dauzère-Pérès, S., Yugma, C., & Wu, C.-H. (2019). Parallel machine scheduling with time constraints on machine qualifications. *Computers & Operations Research*, 107, 61–76.
- [6] Malapert, A., & Nattaf, M. (2019). A new cp-approach for a parallel machine scheduling problem with time constraints on machine qualifications. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 16th International Conference, CPAIOR 2019, Thessaloniki, Greece, June 4–7, 2019, Proceedings 16* (pp. 426–442). Springer.

## The Unreliable Job Selection and Sequencing Problem

Alessandro Agnetis\*

e-mail: [agnetis@unisi.it](mailto:agnetis@unisi.it)

Roel Leus†

e-mail: [roel.leus@kuleuven.be](mailto:roel.leus@kuleuven.be)

Emmeline Perneel†

e-mail: [emmeline.perneel@kuleuven.be](mailto:emmeline.perneel@kuleuven.be)

Ilaria Salvadori\*

e-mail: [ilaria.salvadori@student.unisi.it](mailto:ilaria.salvadori@student.unisi.it)

**Problem statement** We consider a single-machine problem with permanent breakdowns. We are given a set  $J = \{1, \dots, n\}$  of jobs. Each job has a cost  $c_j \geq 0$  which has to be paid if the job is selected, and a reward  $r_j \geq 0$  is achieved if the job is successfully carried out. The machine can fail during the execution of a job and in that event, the job is lost and no further job can be carried out. The probability that the machine fails during the execution of job  $j$  is given by  $1 - \pi_j$ , where  $\pi_j \in [0, 1]$  is called the success probability of job  $j$ . If a certain subset  $S$  of jobs is selected for processing, let  $\sigma$  denote a sequence of these  $|S|$  jobs, and let  $\sigma(k)$  represent the  $k$ -th job in the sequence  $\sigma$ . The probability of reaching and successfully carrying out the  $k$ -th job in the sequence is given by  $\prod_{i=1}^k \pi_{\sigma(i)}$ , therefore the expected reward  $R(S, \sigma)$  from selecting  $S$  and sequencing its jobs according to  $\sigma$  is  $R(S, \sigma) = \sum_{k=1}^{|S|} r_{\sigma(k)} \prod_{i=1}^k \pi_{\sigma(i)}$ . The optimal sequence of jobs to maximize the expected reward is obtained by simply scheduling the jobs by non-increasing values of the following index [2],[3]:

$$Z_j = \frac{\pi_j r_j}{1 - \pi_j}. \quad (1)$$

Consequently, once a subset  $S$  is selected, the jobs in  $S$  will always be sequenced according to a schedule dictated by (1). Letting  $\sigma_Z$  denote such sequence, we let  $R(S) = R(S, \sigma_Z)$ . If we let  $c(S) = \sum_{j \in S} c_j$  denote the total cost of selecting  $S$ , the *expected net profit* of selecting  $S$  is  $z(S) = R(S) - c(S)$ . We study the problem of selecting a subset  $S \subseteq J$  of jobs so that the expected net profit  $z(S)$  is maximized. We call this problem the *Unreliable Job Selection and Sequencing Problem (UJSSP)*.

**Complexity** Two special cases, namely with identical costs and with identical probabilities of success, can both be solved by a greedy algorithm. This greedy algorithm starts with the empty set and iteratively selects a job that yields the largest expected marginal gain in  $z(S)$  until adding an extra job no longer leads to any improvement. This result is deduced from the results from [6] and [1]

---

\*Università di Siena, Dipartimento di Ingegneria dell'Informazione e Scienze Matematiche

†KU Leuven, Research Centre for Operations Research and Statistics

for the case of equal costs and from the results from [5] for the case of equal probabilities. In the general case, we prove that UJSSP is NP-hard by means of a reduction from the *Product Partition Problem (PPP)*, which is known to be strongly NP-hard [4].

**Exact solution methods and computational experiments** We explore several exact solution methods for the problem. Exploiting the ordering induced by (1), we formulate a compact MILP formulation for UJSSP. Using this formulation, all instances with up to 50 jobs are solved to optimality within the 20-minute time limit, whereas no instance with more than 125 jobs is solved within this limit. Secondly, we design a backward dynamic programming (DP) algorithm which can be applied if all costs  $c_j$  are integer. As this approach runs in pseudopolynomial time, we rule out strong NP-hardness for the integer-cost case. The DP algorithm is substantially more scalable than the linear formulation: it solves all instances with up to 2,500 jobs within 10 minutes. However, instances with 3,000 jobs often take more than one hour of computation time. Our main contribution is two stepwise exact methods for solving UJSSP. The stepwise algorithms proceed stepwise through the jobs, forwards or backwards, and iteratively build and prune candidate subsets by exploiting structural properties of the objective function. These algorithms provide the strongest computational performance among all methods considered. When job success probabilities are drawn uniformly all instances are solved in less than one second of CPU time, even for  $n = 10,000$ . For instances with higher success probabilities, computation times reach roughly one minute for the forward algorithm and two minutes for the backward algorithm at  $n = 10,000$ . For these algorithms, a second set of harder instances was constructed from the strongly NP-hard PPP. For these instances, none with more than 50 jobs can be solved to optimality within a 20-minute time limit. Nonetheless, by applying the stepwise methods to instances derived from the PPP, we provide, to the best of our knowledge, the first exact computational results reported for that problem.

## References

- [1] Chade, H., & Smith, L. (2006). Simultaneous search. *Econometrica*, 74(5), 1293-1307.
- [2] Kadane, J.B. (1969). Quiz show problems. *Journal of Mathematical Analysis and Applications*, 27(3), 609-623.
- [3] Mitten, L.G. (1960). An analytic solution to the least cost testing sequence problem. *Journal of Industrial Engineering*, 11(1), 17.
- [4] Ng, C. T., Barketau, M. S., Cheng, T. C. E., & Kovalyov, M. Y. (2010). “Product Partition” and related problems of scheduling and systems reliability: Computational complexity and approximation. *European Journal of Operational Research*, 207(2), 601-604.
- [5] Olszewski, W., & Vohra, R. (2016). Simultaneous selection. *Discrete Applied Mathematics*, 200, 161-169.
- [6] Stadje, W. (1995). Selecting jobs for scheduling on a machine subject to failure. *Discrete Applied Mathematics*, 65(3), 257-265.

# A comprehensive analysis of the sparrow-3D heuristic: Bottlenecks in Irregular Strip Packing

Jan Martens

KU Leuven, NUMA

e-mail: [jan.martens@kuleuven.be](mailto:jan.martens@kuleuven.be)

Jonas Tolleneare

KU Leuven, NUMA

e-mail: [jonas.tolleneare@kuleuven.be](mailto:jonas.tolleneare@kuleuven.be)

January 29, 2026

The 3D irregular strip packing problem (or 3D nesting problem) is a cutting and packing problem in which 3D irregular objects need to be positioned within a 3D cuboid container with fixed width and depth, such that the total height of the container is minimized. This challenging optimization problem has several practical applications in additive manufacturing and logistics.

The recent adaptation of the sparrow algorithm (a heuristic for 2D nesting problems [1]) to the 3D irregular strip packing problem has yielded state-of-the-art results. However, while the sparrow-3D algorithm performs effectively on standard benchmark instances, computational efficiency remains a significant challenge when scaling to complex instances with high-resolution geometry and a wider variety of scales between items. Our current research shifts the focus from prototype implementation to diagnosis, presenting a thorough analysis of the heuristic.

Through extensive profiling and experimental validation on new and existing benchmark instances, we identify that collision detection and quantification constitute the primary computational bottleneck. To mitigate this, we argue that a fundamental restructuring of the collision detection engine is required. Additionally, we analyse and categorize cases in which the heuristic provides insufficient guidance, causing the search to stagnate in local minima. Based on this analysis, we propose a strategy for integrating a more efficient collision detection engine and refining the heuristic guidance to escape local optima.

We conclude by outlining our ongoing efforts to develop an enhanced variant of the heuristic, aiming to retain accuracy while speeding up the search process and reducing computational load.

## References

- [1] Jeroen Gardeyn, Greet Vanden Berghe, Tony Wauters, 2025, An Open-Source Heuristic to Reboot 2D Nesting Research.

# Efficient Compliance Checking of EU Driver Regulations within Local Search-Based VRP Solvers

Niels De Walsche

KU Leuven, Numerical Analysis and Applied Mathematics (NUMA)

e-mail: [niels.dewalsche@kuleuven.be](mailto:niels.dewalsche@kuleuven.be)

Employee wellbeing has received increased attention and prioritization by employers in recent years. This is especially true in the transportation sector, where the European Union has introduced regulations to ensure that drivers are not overworked.

The Truck Driver Scheduling Problem (TDSP), introduced by Goel (2009), aims to create schedules that comply with the aforementioned European Union regulations. Ensuring compliant schedules is a crucial component of the Vehicle Routing and Truck Driver Scheduling Problem (VRTDSP), where the goal is to find routes for the trucks and schedules for the drivers that minimize both the number of drivers as well as the total distance traveled.

However, determining the feasibility of a route remains a difficult problem. This is especially true in a local search context, where many routes are explored and thus need to be checked for feasibility. To address this challenge, we propose an efficient algorithm for checking the compliance of truck driver schedules with EU regulations. This algorithm is designed to be integrated into local search-based VRP solvers, enabling rapid feasibility checks during the optimization process. By exploiting the properties of the local search framework, our algorithm significantly reduces the computational overhead associated with compliance checking, thereby efficiently solving the VRTDSP.

We conduct experiments with the proposed algorithm on a set of benchmark instances from the literature. The results indicate the effectiveness of our approach, with it outperforming prior heuristic approaches in terms of both solution quality and computation time.

# An order-first split-second framework for solving the 2D guillotine cutting stock problem with bounded-sized stock

Khadija Hadj Salem

KU Leuven, Department of Computer Science, NUMA, Gent, Belgium  
e-mail: [khadijaalice.hadjsalem@kuleuven.be](mailto:khadijaalice.hadjsalem@kuleuven.be)

Tony Wauters

KU Leuven, Department of Computer Science, NUMA, Gent, Belgium  
e-mail: [tony.wauters@kuleuven.be](mailto:tony.wauters@kuleuven.be)

Cutting stock problems are challenging combinatorial optimization problems whenever stock materials need to be cut into smaller pieces. These problems arise in many industries, such as the glass manufacturing industry (see [4], steel (see [3]), textiles (see [2], and the honeycomb cardboard industry (see [5]), among others.

In this work, we consider the two-dimensional guillotine bounded-sized cutting stock problem (2D-BSCSP), which was first introduced in [1]. According to the typology in [6], the 2D-BSCSP is a new variant of the classical 2D-CSP with predefined and fixed dimensions, which is proven to be a strongly  $\mathcal{NP}$ -hard problem. The 2D-BSCSP involves cutting a set of small rectangular pieces from a large plate whose dimensions vary within a specific range. This variability introduces additional complexity, as the plate dimensions are not fixed and must be considered when optimizing the cutting process, which involves two non-exact guillotine stages.

We propose a general framework based on a binary sequence of decisions. Each decision corresponds to an exploration strategy on a graph. The resulting path defines the construction of a sequence of strips. This strip sequence is then optimally packed into cutting patterns using a polynomial-time exact method. Multiple strip sequences may lead to the same cutting solution. These equivalent representations are referred to as symmetries. To address this issue, we introduce a sorting bias that reduces symmetry and guides the search toward canonical representations. Using the evaluation strategy described above, a metaheuristic can focus its exploration on the binary decision sequence, while the quality of each solution is assessed exactly and efficiently. To illustrate the proposed framework, we present implementations based on a Genetic Algorithm and Simulated Annealing.

The proposed framework offers two main advantages over traditional approaches that rely on generating and selecting cutting patterns. First, it provides direct access to an exact polynomial evaluation method, ensuring accurate fitness assessment. Second, its simple binary representation naturally enables the appli-

cation of a wide range of metaheuristics, including local search methods that are rarely used to solve cutting and packing problems.

**Keywords:** Cutting stock problem; Bounded-sized; 2-stage guillotine; Genetic Algorithm; Simulated Annealing; Split-and-Sort.

## References

- [1] Khadija Hadj Salem and Alexis Robbes. A large neighborhood search matheuristic for the two-dimensional guillotine bounded-sized cutting stock problem. *European Journal of Operational Research*, 2025. work under review since August 5.
- [2] Khadija Hadj Salem, Elsa Silva, José Fernando Oliveira, and Maria Antónia Carravilla. Mathematical models for the two-dimensional variable-sized cutting stock problem in the home textile industry. *European Journal of Operational Research*, 306(2):549–566, 2023.
- [3] Yunfeng Ma, Jiayi Zhang, Xijie Yang, Jihao Li, Xiaoxin Su, and Haoxun Chen. Two-dimensional cutting stock problem with flexible length and usable leftovers in the steel industry. *European Journal of Operational Research*, 326(2):207–219, 2025.
- [4] F. Parreño, M.T. Alonso, and R. Alvarez-Valdes. Solving a large cutting problem in the glass manufacturing industry. *European Journal of Operational Research*, 287(1):378–388, 2020.
- [5] Paula Terán-Viadero, Antonio Alonso-Ayuso, and F. Javier Martín-Campo. Mathematical optimisation in the honeycomb cardboard industry: A model for the two-dimensional variable-sized cutting stock problem. *European Journal of Operational Research*, 319(1):303–315, 2024.
- [6] Gerhard Wäscher, Heike Haußner, and Holger Schumann. An improved typology of cutting and packing problems. *European Journal of Operational Research*, 183(3):1109–1130, 2007.

# Analysis of the relationship between different constraint types in personnel rostering problems

Lisa Garcia Tercero

KU Leuven, Department of Computer Science

UGent, Faculty of Economics and Business Administration

e-mail: [lisa.garciatercero@kuleuven.be](mailto:lisa.garciatercero@kuleuven.be)

Personnel rostering problems are often subject to a lot of constraints. While some are enforced by law or by the employee's contract (e.g., the number of hours that should be worked per week), others are added by the employer to make the roster more desirable for the employees (e.g., by imposing a maximum on the number of consecutive days worked or by taking into account employees' preferences). When the problem contains a lot of soft constraints, it is hard to evaluate a roster in the end since there are so many objectives to consider. This is especially challenging when planning tools are being used by planners in the company who do not know the underlying model itself. The only option they have to alter a roster is often to change the values of some parameters, of which there are many and with many options to choose from. Changing these parameters the wrong way may lead to completely useless rosters, even though they are mathematically optimal.

In this work we investigate the impact of the parameters related to hard constraints on the value of soft constraints. Specifically, we compute lower and upper bounds on the values an objective function can reach given the chosen parameters for the hard constraints. If these bounds are close enough to each other, we can argue that this objective will always fall within an acceptable range and can therefore be removed from the model, thus reducing complexity.

We start by considering two types of constraints. *Counter constraints* enforce how many units of time an employee can minimally and/or maximally work or have off during a given period of time. These units of time can be hours, shifts, days, weekends, etc. On the other hand, *series constraints* impose limits on series of consecutive time units worked or idle. We consider the series constraints to be hard and investigate their impact on the counter constraints. Our aim is to answer the question: "What are the minimum and maximum units of time an employee can work during a given period, given the series constraints?" To illustrate this, consider the unit of time to be a day. Additional to the length of the period, the model includes four parameters for series constraints:

- $d$  = number of days in the time horizon
- $days\_worked\_min$  = minimum number of consecutive days worked
- $days\_worked\_max$  = maximum number of consecutive days worked

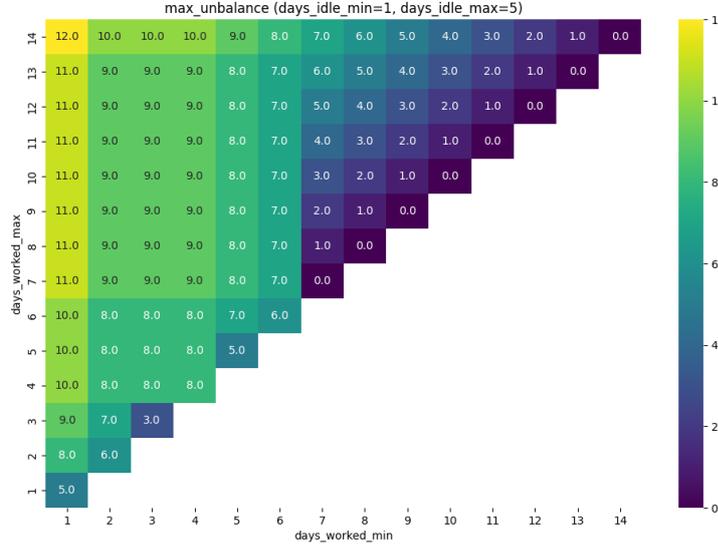


Figure 1: The value of  $D_{max} - D_{min}$  in function of  $days\_worked\_min$  and  $days\_worked\_max$  when  $days\_idle\_min$  and  $days\_idle\_max$  are fixed.

- $days\_idle\_min$  = minimum number of consecutive days off
- $days\_idle\_max$  = maximum number of consecutive days off

We aim to compute the maximum number of days an employee can work in the given period of time ( $= D_{max}$ ) as well as the minimum number of days worked ( $= D_{min}$ ). If these extreme values (or the difference between them) are already deemed acceptable, we can omit them from the model and focus on optimizing other objectives.

The values of  $D_{min}$  and  $D_{max}$  are computed using analytical insights in the problem and verified with an IP-model. By means of example, we show how the difference between  $D_{min}$  and  $D_{max}$  varies when  $days\_worked\_min$  and  $days\_worked\_max$  vary in a period of  $d = 14$  days. The values of  $days\_idle\_min$  and  $days\_idle\_max$  are fixed to 1 and 5, respectively. The results are shown in Figure 1. They show that the relationship between these values is not linear. For example, when  $days\_worked\_min \in [2, 4]$  the  $days\_worked\_max$  parameter can range from 7 to 13 without impacting bounds on the counter constraints. This small example already provides useful insights. Including more hard constraints, such as constraints on weekends, will further decrease the value of  $D_{max} - D_{min}$  and thereby control the values of counter constraints without explicitly having to optimize them in a model.

# Special cases of the Joint Replenishment Problem (JRP): A mathematical analysis

Tim Stephan

UCLouvain, LIDAM/CORE

e-mail: [tim.stephan@uclouvain.be](mailto:tim.stephan@uclouvain.be)

Stefan Creemers

UCLouvain, LIDAM/CORE

e-mail: [stefan.creemers@uclouvain.be](mailto:stefan.creemers@uclouvain.be)

December 5, 2025

The Joint Replenishment Problem (JRP) is a classical inventory control problem. But even for small instances, optimal policies are poorly understood: available structural results are limited, exact evaluation typically relies on Markov chains or simulation, and no closed-form expressions for optimal long-run costs are known [1, 2]. To gain a deeper theoretical understanding of the JRP, this paper investigates three fundamental special cases obtained by eliminating one cost component: no holding costs ( $h = 0$ ), no major order costs ( $K = 0$ ), and no minor order costs ( $k = 0$ ).

First, we show that the case  $h = 0$  is straightforward: items with zero holding cost can be removed from the problem in its formulation without limitations on inventory capacity or order sizes, as their contribution to long-run average costs vanishes.

Assuming  $h > 0$  and restricting to the space of deterministic policies without throwing away inventory, we show the existence of optimal policies in a further restricted space of policies placing orders if and only if an item with no leftover quantity is demanded.

For  $K = 0$ , the JRP decomposes into independent single-item problems. We provide a complete mathematical characterization of optimal policies for both finite-horizon and long-run average cost criteria. Our analysis reveals several structural insights that do not always align with classical Economic Order Quantity (EOQ) intuition. Notably, the long-run EOQ differs from the finite-horizon EOQ with convergence in the limit.

For  $k = 0$ , we prove the existence of optimal policies with unique reorder levels and derive for the first time closed-form expressions for the expected long-run costs. These expressions allow explicit optimization of reorder levels and offer structural insights that can be leveraged to accelerate computation of optimal JRP policies.

Finally, we present several counterexamples illustrating subtle phenomena that arise even in simplified JRPs.

## References

- [1] Stefan Creemers and Robert Boute. The joint replenishment problem: Optimal policy and exact evaluation method. *European Journal of Operational Research*, 302(3):1175–1188, 2022.
- [2] Edson Antônio Gonçalves de Souza, Tanja Mlinar, Maud Van den Broeke, and Stefan Creemers. Dynamic programming in inventory management: A review. *Computers & Operations Research*, 183:107164, 2025.

# Stochastic bilevel optimization with dynamic decision-making followers

Diego Jiménez

SKEMA Business School/KU Leuven, Faculty of Economics and Business

e-mail: [diego.jimenez@kuleuven.be](mailto:diego.jimenez@kuleuven.be)

Bernardo K. Pagnoncelli

SKEMA Business School

e-mail: [benardo.pagnoncelli@skema.edu](mailto:benardo.pagnoncelli@skema.edu)

Hande Yaman

KU Leuven, Faculty of Economics and Business

e-mail: [hande.yaman@kuleuven.be](mailto:hande.yaman@kuleuven.be)

We study a hierarchical setting where a leader interacts with multiple independent followers that takes *sequential discrete decisions* over time, within a stochastic environment. The leader can be interpreted as a coordinator who influences followers' behavior through discrete discounts, thereby minimizing its own cost function:

$$c(\mathbf{x}) = f(\mathbf{x}) + \rho[g(\mathcal{P}(\mathbf{x}))]. \quad (1)$$

Here,  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{B}^{n_x}$  is the *vector of discounts* offered by the leader, and  $\mathcal{P}(\mathbf{x})$  is the collection of *followers' policies*, representing their potential behavior within the stochastic environment, given the leader's discounts vector  $\mathbf{x}$ . The cost function (1) is composed then by two terms: (i) the direct cost incurred when giving discounts  $f(\mathbf{x})$ , and (ii) a risk measure  $\rho$  of the cost  $g$ , which depends on followers' policies  $\mathcal{P}(\mathbf{x})$ .

Consequently, the leader's trade-off is in between how much discounts to give to avoid possible followers' undesired behavior. For example, in inventory management (IM), followers place orders to a single supplier, who seeks to provide price discounts so that its delivery capacity is never exceeded. Consider then the set of followers  $\mathcal{F}$ , and let  $\pi_f^{\mathbf{x}}$  a feasible policy of follower  $f$ , with  $f \in \mathcal{F}$ . Defining  $\mathcal{P}(\mathbf{x}) = \{\pi_f^{\mathbf{x}}\}_{f \in \mathcal{F}}$ , our target problem can be written as:

$$\min \left\{ c(\mathbf{x}) : \mathbf{x} \in \mathcal{X}, \pi_f^{\mathbf{x}} \in \Pi_f^*(\mathbf{x}), \forall f \in \mathcal{F} \right\} \quad (2)$$

where  $\Pi_f^*(\mathbf{x})$  is the set of optimal policies for follower  $f$ , given leader's discounts  $\mathbf{x}$ , expressing the fact that policies  $\pi_f^{\mathbf{x}}$  must be optimal.

On the other hand, followers are assumed to be following an Infinite Horizon Markov Decision Process (IHMDP) [1], a suitable framework to model discrete dynamic decision-making under uncertainty. Using standard MDP notation, each

follower is characterized by *discrete* state and action spaces,  $\mathcal{S}_f$  and  $\mathcal{A}_f$ , respectively.  $\mathbb{P}_f(s'|s, a)$  is the probability of transition from state  $s$  to  $s'$ , when performing action  $a$ , and an instant cost  $r_{s,a}^f(\mathbf{x})$  is incurred by follower  $f$ , when taking action  $a$  at state  $s$ . A *policy*  $\pi_f^{\mathbf{x}} : \mathcal{S}_f \rightarrow \mathcal{A}_f$  for follower  $f$  is then a mapping that assigns an unique action  $a \in \mathcal{A}_f$  to each state  $s \in \mathcal{S}_f$ . An optimal policy minimizes the value function  $v_s^f$ , for every  $s \in \mathcal{S}_f$ :

$$v_s^f = \min_{\pi_f^{\mathbf{x}}(s) \in \mathcal{A}_f} \left\{ r_{s, \pi_f^{\mathbf{x}}(s)}^f(\mathbf{x}) + \lambda_f \sum_{s' \in \mathcal{S}_f} \mathbb{P}_f(s'|s, \pi_f^{\mathbf{x}}(s)) v_{s'}^f \right\}, \quad (3)$$

where  $\lambda_f$  is the discount factor for follower  $f$ . In this work, we study the case where certain combinations of followers' actions are costly for the leader, specifically when their total action exceeds a given threshold  $K$ . To do so, a binary representation of followers' policies  $\pi_{s,a}^f \in \{0, 1\}$  is used, where value 1 indicates that follower  $f$  takes action  $a$ , at state  $s$ . Thus, we define the leader's deviation cost:

$$g(\mathcal{P}(\mathbf{x})) = c_{pen} \cdot \max \left\{ \sum_{f \in \mathcal{F}} \sum_{s \in \mathcal{S}_f} a \cdot \pi_{s,a}^f - K, 0 \right\}, \quad (4)$$

which represents the cost of the total action positive deviation from  $K$ . Therefore, we propose a mixed-integer linear programming (MILP) formulation for solving problem (2), for the worst-case value of (4). We consider that both  $f(\mathbf{x})$  and  $r_{s,a}^f(\mathbf{x})$  are linear in  $\mathbf{x}$ , giving place to our base formulation:

$$\min_{\theta_f, \eta \geq 0} f(\mathbf{x}) + c_{pen} \cdot \eta \quad (5a)$$

$$\text{s.t. } \mathbf{x} \in \mathcal{X} \quad (5b)$$

$$\eta \geq \sum_{f \in \mathcal{F}} \theta_f - K \quad (5c)$$

$$\theta_f \geq \sum_{a \in \mathcal{A}_f} a \cdot \pi_{s,a}^f \quad s \in \mathcal{S}_f, f \in \mathcal{F} \quad (5d)$$

$$\{\pi_{s,a}^f\}_{s \in \mathcal{S}_f, a \in \mathcal{A}_f} \in \Pi_f^*(\mathbf{x}) \quad f \in \mathcal{F} \quad (5e)$$

where the description of the optimal set of policies  $\Pi_f^*(\mathbf{x})$  is based on the Karush-Kuhn-Tucker (KKT) conditions. Preliminary experiments on IM instances show that regular MILP solvers fail to solve formulation (5) to optimality within a time limit of 300s, even for very modest state and action spaces. This is due to the weakness of the linear relaxation of formulation (5), particularly because of the big-M description of  $\Pi_f^*(\mathbf{x})$  [2]. To this end, we leverage the strong duality condition to obtain a tighter formulation. Experiments show significant improvements in computational performance.

## References

- [1] M. L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming. Hoboken, New Jersey, U.S.: Wiley, 1994.
- [2] T. Kleinert, M. Labbé, F. Plein, and M. Schmidt, "Technical note—there's no free lunch: On the hardness of choosing a correct big-M in bilevel optimization," *Operations Research*, vol. 68, no. 6, pp. 1716–1721, 2020.

# Adaptive Hybrid VNS-LNS for Time-Dependent Grocery Delivery Vehicle Routing

Omar Elsherif<sup>c,\*</sup>

e-mail: [omar.mohammedelsherif@universitietantwerpen.be](mailto:omar.mohammedelsherif@universitietantwerpen.be)

Kenneth Sörensen\*      Bryan Galarza Montenegro\*  
An Caris\*\*

<sup>c</sup> Corresponding author.

\* Antwerp Operations Research Group - ANT/OR, Department of Engineering Management (ENM), University of Antwerp, Belgium.

\*\* Research Group Logistics, Hasselt University, Belgium.

Grocery logistics present unique operational challenges. Perishable products (i.e. groceries) require multi-compartment vehicles, narrow delivery time windows, and complex urban traffic patterns. We address a multi-depot, multi-compartment vehicle routing problem with time-dependent travel times and time windows (TD-VRPTW), incorporating realistic constraints related to grocery delivery: vehicles respect compartment specific capacity limits, deliver within customer time windows, and account for traffic congestion across planning periods (morning, midday, evening). The objective is to minimize total operational costs including fixed vehicle deployment and distance dependent variable costs.

We develop a hybrid metaheuristic implemented in Julia, integrating Variable Neighborhood Search (VNS) with Large Neighborhood Search (LNS) through adaptive operator selection. The algorithm workflow (Figure 1) consists of four phases: The first phase is **Initialization** that constructs an initial solution using regret-3 insertion, prioritizing customers with highest insertion regret. The second phase is **VNS Local Search** that applies three operators sequentially (*relocate*, *swap*, *2-opt*) with candidate lists of neighborhood sizes and first-improvement acceptance. The third phase is **LNS Diversification** (Mancini, 2016) that removes 10-30% of customers using four destroy operators (*random*, *worst*, *related*, *shaw*) and reinserts them via three repair operators (*greedy*, *regret-2*, *regret-3*). The final phase is **Adaptive Selection** for updating operator weights using ALNS-style scoring as stated by Ropke and Pisinger (2006):  $w_{\text{new}} = 0.9 \cdot w_{\text{old}} + 0.1 \cdot \sigma$ , where  $\sigma \in \{50, 33, 13\}$  rewards solution improvements. The algorithm alternates VNS-VNS-LNS in a 2:1 ratio, balancing local refinement with diversification. Implementation uses piecewise travel time calculation across traffic periods and maintains multi-compartment load tracking with feasibility checking at each move.

We tested the solution methodology on problem instances in Antwerp. These experiments contained 10-50 customers, 2 depots, 2-3 heterogeneous vehicles, and had realistic time-dependent travel times (obtained from OpenRouteService API). The algorithm matches the commercial solver Gurobi 10.0, obtaining optimal solutions for vehicle routing costs across four scenarios with different

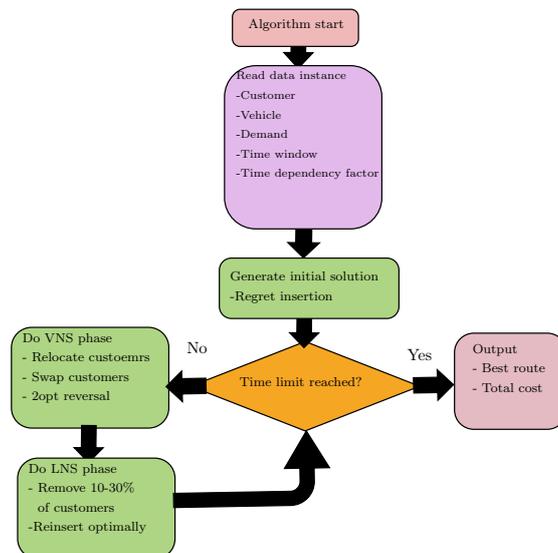


Figure 1: Flowchart of Adaptive Hybrid VNS/LNS for Time-Dependent Grocery Delivery VRP. The algorithm alternates between VNS local search and LNS diversification phases with adaptive operator selection based on historical performance.

settings: single-period, two-period (flexible and fixed), and three-period TD-VRPTW. Compared to **Gurobi**, our approach achieved the following results:

1. **Single-period:** Total vehicle routing cost €89.69, 5s vs. 155s (96% faster) with optimality gap 0%.
2. **Two-period flexible:** Total vehicle routing cost €90.32, 6.2s vs. 1960s (99.7% faster) with optimality gap 0%.
3. **Two-period fixed:** Total vehicle routing cost €104.63, 2.2s vs. 2500s (99.9% faster) with optimality gap 0%.
4. **Three-period:** Total vehicle routing cost €104.37, 10s vs. 2800s (99.6% faster) with optimality gap 0%.

This research demonstrates that adaptive hybrid metaheuristics with ALNS-style operator weighting achieves robust performance across varying conditions while reducing computation time by over 96%, making it suitable for real-time operational planning in grocery logistics.

## References

- Mancini, S. (2016). A real-life multi depot multi period vehicle routing problem with a heterogeneous fleet: Formulation and adaptive large neighborhood search based matheuristic. *Transportation Research Part C: Emerging Technologies*, 70:100–112.
- Ropke, S. and Pisinger, D. (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation science*, 40(4):455–472.

# Contextual optimization with recourse for production-transshipment problem with uncertain yield and demand

Shi Zong

KU Leuven, Institute for Mobility  
e-mail: [shi.zong@kuleuven.be](mailto:shi.zong@kuleuven.be)

Johan Joubert

KU Leuven, Institute for Mobility  
e-mail: [johan.joubert@kuleuven.be](mailto:johan.joubert@kuleuven.be)

Uncertain yield and demand exacerbate the complexity of supply chain decision-making, undermining overall supply chain efficiency. Prevailing approaches often rely on deterministic assumptions or unconditional distributions to represent uncertainty, potentially oversimplifying complex dynamics in real-world scenarios. In this study, we investigate a production-transshipment problem with uncertain yield and demand within a two-echelon supply chain structure, consisting of production, transshipment, and demand stages. Given this setting as a representative case, we address two primary research questions. First, how can supply chain uncertainty be addressed more effectively, efficiently and reliably using modern data resources? Second, while most existing approaches assume deterministic constraints, how should supply chain models be addressed when uncertain parameters, such as demand and yield, are inherently embedded within the constraints [3]?

Traditional approaches generally fall into two categories. The "predict-then-optimize" (PTO) framework uses machine learning to predict uncertain parameter values, without considering how those predictions affect the quality of decisions, often leading to suboptimal or even wrong decisions when actual yield or demand deviates [2]. Stochastic programming (e.g., SAA) accounts for distributional information but seeks a one-size-fits-all solution that performs well on average across all scenarios, failing to leverage contextual information that could better discriminate among them [1].

To overcome these limitations, we propose contextual optimization with recourse. We employ machine learning methods (e.g., k-nearest neighbors, random forest) not to directly predict parameter values, but rather to generate weight vectors that measure the similarity between the current operational context (such as the date or specific machine states) and historical scenarios. This enables the model to prioritize scenarios most relevant to the current situation, effectively leveraging contextual information from modern data resources. A critical innovation of our approach is addressing the challenge where uncertain yield and demand are embedded directly within the flow balance constraints. Instead of enforcing rigid constraints based on static predictions, we incorporate the recourse mechanism from two-stage stochastic programming by introducing "wait-and-see" recourse variables that absorb deviations between planned flows and realized conditions. By minimizing the sum of transportation costs and expected recourse costs across contextually weighted scenarios, our model adjusts production and

transshipment flows to balance efficiency with the feasibility risks unique to the current environment.

We implemented our proposed contextual optimization with recourse by integrating two machine learning methods, k-nearest neighbors and random forest, and compared their performance with two benchmark methods (PTO and SAA). Experiments across multiple problem scales demonstrate that the proposed approach achieves an average cost reduction of 18.32% compared to SAA and 33.56% compared to PTO. The results indicate that by leveraging contextual information to weight historical scenarios, the model effectively manages uncertain parameters embedded in constraints, offering an effective solution for complex supply chain systems.

## References

- [1] Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3), 1025-1044.
- [2] Elmachtoub, A. N., & Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1), 9-26.
- [3] Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., & Vidal, T. (2025). A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2), 271-289.

# From Waste to Worth: Extending The Multi-Echelon Network Design Problem Towards Circular Supply Chain Modelling

Alix Langenaeker

University of Antwerp, Department of Engineering Management  
Flanders Make@UAntwerp  
e-mail: [alix.langenaeker@uantwerpen.be](mailto:alix.langenaeker@uantwerpen.be)

Philippe Nimmegeers

University of Antwerp, Department of Engineering Management  
Flanders Make@UAntwerp, NANOligh Centre of Excellence  
e-mail: [philippe.nimmegeers@uantwerpen.be](mailto:philippe.nimmegeers@uantwerpen.be)

Christof Defryn

University of Antwerp, Department of Engineering Management  
e-mail: [christof.defryn@uantwerpen.be](mailto:christof.defryn@uantwerpen.be)

The rising demand for food, energy, water, and raw materials, coupled with population growth, climate change, and resource scarcity, underscores the urgency of developing smarter and sustainable systems. Circular supply chains play a crucial role in meeting these challenges by improving efficiency, lowering costs, and strengthening environmental resilience. Circularity aims to maintain the value of resources at their highest functionality, minimizing waste through strategies as those outlined in, e.g., the 9R framework (Refuse, Rethink, Reduce, Reuse, Repair, Refurbish, Remanufacture, Repurpose, Recycle, Recover) [1, 3]. However, the optimization of sustainable circular supply chains remains challenging, largely due to the absence of a generic quantification framework to evaluate and enhance circularity and sustainability at the same time.

Transitioning from linear to circular supply chains introduces new processes and dynamics. From a high-level perspective, two new echelons emerge which should be established and integrated within the supply chain network: (1) waste collection and (2) reprocessing (covering various circular strategies and technologies).

This study addresses the pressing need for a generic guiding framework by proposing an optimisation model, applicable to different industries and supply chains through case-specific data. The critical challenge lies in balancing the volume and complexity of the required data with the ability to capture the case-specific characteristics, essential for circular supply chain design.

The created model builds on a multi-echelon network design problem, providing a solid basis for strategic decisions such as the optimal location and capacity planning. To tailor the multi-echelon network design problem towards the con-

text of a circular supply network, various assumptions must be revisited [2]. This study proposes a set of extensions to the traditional multi-echelon network design problem to capture these specificities.

The contribution of this study is threefold. In a first extension, the assessment and selection of the most appropriate technology is introduced as a strategic decision, both for the collection process type and the reprocessing facility technology. This technology decision is strongly connected to the different R-strategies. Next, the introduction of technology decisions is supplemented by an important and realistic constraint in circular supply chain design on compatibility between heterogeneous waste streams and, respectively, the set of collection types and reprocessing technology options. Finally, the assumption of fixed waste quality is reconsidered, introducing a potential waste quality degradation related to the selection of a certain collection strategy.

The proposed framework outlines a more realistic representation of the network design problem in the circular supply chain setting, offering a strong basis for constructing generic optimization models to support strategic decision-making.

## References

- [1] Grow Circular. (n.d.). "9R Framework". <https://grow-circular.eu/knowledge-base/9r-framework/> (accessed: October 2, 2024)
- [2] Liao, T.-Y. (2018). "Reverse logistics network design for product recovery and remanufacturing". *Applied Mathematical Modelling*, 60, 145-163. <https://doi.org/10.1016/j.apm.2018.03.003>
- [3] Potocnik, J. (2013). "Towards the circular economy—economic and business rationale for an accelerated transition". *Ellen MacArthur Foundation*.

# Decision support for afforestation planning: Integrating exact and heuristic spatial optimization methods

Grethell Castillo-Reyes

University of Informatics Sciences, Center for Computational Mathematics Studies

e-mail: `gcreyes@uci.cu`

Dirk Roose

KU Leuven, Department of Computer Science

e-mail: `dirk.roose@kuleuven.be`

Greet Vanden Berghe

KU Leuven, Department of Computer Science

e-mail: `greet.vanden.berghe@kuleuven.be`

Land-use planning often requires spatial optimization methods that account for interactions between neighboring locations. Afforestation of a part of a river catchment can prevent erosion and thus the amount of sediment that leaves the catchment. The simulation of sediment flow is based on information about the sediment produced at each site in the catchment and on models that describe how and to what extent the sediment flows to neighboring sites. The optimization problem involves determining which sites should be afforested to minimize sediment leaving the catchment, subject to a budget constraint, such as the size of the afforested area.

We develop an Integer Programming (IP) formulation that integrates raster-based information on elevation, sediment production, a flow routing model, transport capacities, river locations, and candidate afforestation sites (i.e., sites eligible for afforestation). The optimization model assigns to each candidate cell a binary decision variable, subject to a global budget constraint. For every cell, the model computes incoming sediment and outgoing sediment to its downstream neighbors. The outgoing sediment is constrained by a piecewise concave linear transport function based on transport capacity, reflecting hydrological processes. The objective is to minimize the incoming sediment at a designated target cell. The implementation in Gurobi outputs optimal afforestation locations, estimated sediment reductions, and geospatial raster files for direct integration into Geographic Information Systems. This work is a significant improvement on the IP formulation for the afforestation problem in [1], which used a simple flow routing model and a convex piecewise linear transport function, while its implementation in Lingo proved to be effective only for very small datasets.

We compare the results of the IP formulation with those obtained using CAMF, a heuristic for solving this afforestation optimization problem based on

the steepest ascent hill-climbing algorithm, which generates good quality solutions while iteratively capturing spatial interactions. CAMF was developed in [1] and extended in several directions, including the incorporation of several flow network models and parallelism for multi-core computers in [2]. CAMF has been used to solve the afforestation problem for large catchments (with up to 1 000 000 cells) and for a large number of cells to be afforested.

The experiments use a subcatchment of the Tabacay river catchment in Ecuador, represented by 15 583 raster cells, of which 10 675 cells are identified as candidate cells for afforestation. Computations were performed on a 14-core workstation. Figure 1 compares the cells selected by the IP model and CAMF for up to 200 afforestation cells, showing their overlap. Although CAMF does not always choose exactly the same cells as the optimal IP solution, it generally shows strong spatial agreement with similar sediment loss reduction. The proposed IP model offers a consistent accuracy and scalability up to medium-sized selections, reaffirming its value as a benchmark for evaluating the heuristic method and guiding further improvements in spatial decision-making tools. Additional tests and improvements for larger problem sizes are needed.

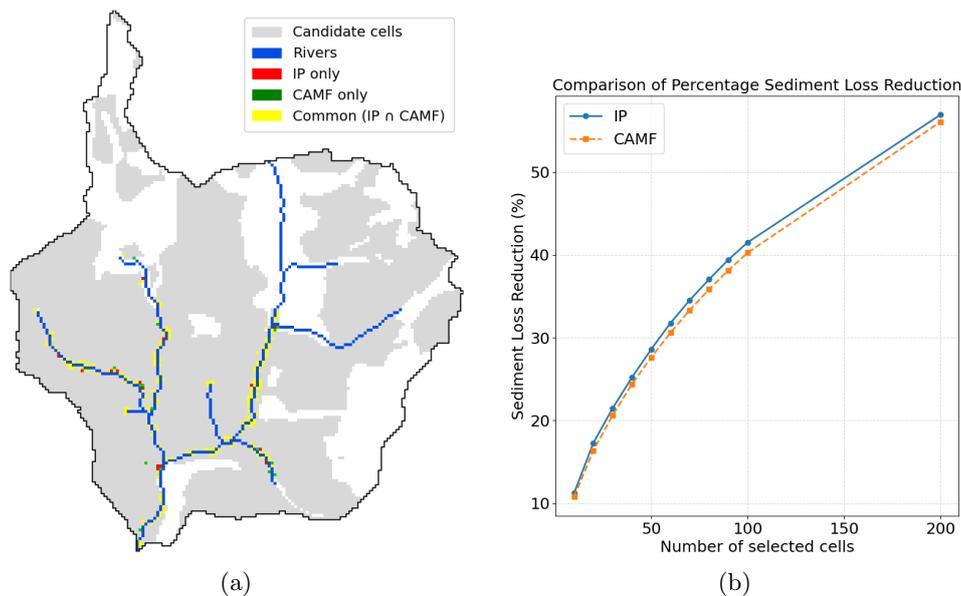


Figure 1: Comparison between IP model and CAMF heuristic: (a) common and unique cell selection for each approach; (b) comparison of sediment loss reduction.

## References

- [1] Vanegas, P., Cattrysse, D., & Van Orshoven, J. (2012). Allocating reforestation areas for sediment flow minimization: An integer programming formulation and a heuristic solution method. *Optimization and Engineering*, 13, 247–269.
- [2] Castillo-Reyes, G. (2024). *Extension and acceleration of the CAMF approach to select optimal intervention sites in the presence of spatial interaction* (PhD thesis). KU Leuven and University of Informatics Sciences.

# Sustainability Analytics for Assessing Environmental Reporting and Media Perception

Fernando Diaz Gonzalez Manjarrez

IÉSEG School of Management, e-mail: [f.diazgonzalez@ieseg.fr](mailto:f.diazgonzalez@ieseg.fr)

Lies Bouten

IÉSEG School of Management, e-mail: [l.bouten@ieseg.fr](mailto:l.bouten@ieseg.fr)

Matthijs Meire

IÉSEG School of Management, e-mail: [m.meire@ieseg.fr](mailto:m.meire@ieseg.fr)

Kristof Coussement

IÉSEG School of Management, e-mail: [k.coussement@ieseg.fr](mailto:k.coussement@ieseg.fr)

Sustainability analytics refers to the systematic use of technology for gathering, sharing, examining, and applying information related to sustainability throughout an organization [9]. According to [7], incorporating environmental, health, and safety (EHS), as well as sustainability analytics, enables companies to minimize risks and achieve cost efficiencies. Recent advances in natural language processing (NLP) and machine learning have played growing role in sustainability analytics, such as in assessing corporate sustainable disclosures [4, 10, 6]. We contribute to this line of work by analyzing sustainability reports of 89 companies belonging to the Euronext 100 index from 2018 to 2022, to know how companies' varying levels of engagement with the Task Force on Climate-related Financial Disclosures (TCFD) sustainability framework, affect the news media opinion of the companies. Based on [6] we used text analytics metrics such as *tf\_idf* to identify several levels of engagement of the company with the TCFD. Moreover, we use climate change-related news, filtered by ClimateBERT [10, 4] and quantified their opinion with the Janis-Fadner coefficient of imbalance [8, 1, 2]. We further explore the moderating roles of a company's environmental footprint and its classification as an environmentally sensitive industry [5, 3, 4]. Preliminary results show that TCFD engagement levels exhibit distinct associations with news media opinion. Higher environmental footprint generally worsens opinion, but specific TCFD engagement levels offset this effect. In environmentally sensitive industries, past news media opinion is more predictive of current news media opinion than their TCFD levels of engagement. Future work will focus on predictive modeling on company value, optimization of disclosure strategies, and the development of NLP based tools for automated sustainability reports assessment.

## References

- [1] Walter Aerts and Denis Cormier. Media legitimacy and corporate environmental communication. *Accounting, Organizations and Society*, 34(1):1–27, 1 2009.
- [2] Pratima Bansal and Iain Clelland. Talking Trash: Legitimacy, Impression Management, and Unsystematic Risk in the Context of the Natural Environment. *The Academy of Management Journal*, 47(1):93–103, 2004.
- [3] Pascual Berrone, Andrea Fosfuri, and Liliana Gelabert. Does Greenwashing Pay Off? Understanding the Relationship Between Environmental actions and environmental legitimacy. *Journal of Business Ethics*, 144(2):363–379, 2017.
- [4] Julia Anna Bingler, Mathias Kraus, Markus Leippold, and Nicolas Webersinke. How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. *Journal of Banking and Finance*, 164, 7 2024.
- [5] Charles H. Cho and Dennis M. Patten. The role of environmental disclosures as tools of legitimacy: A research note. *Accounting, Organizations and Society*, 32(7-8):639–647, 10 2007.
- [6] Dong Ding, Bin Liu, and Millicent Chang. Carbon Emissions and TCFD Aligned Climate-Related Information Disclosures. *Journal of Business Ethics*, 182(4):967–1001, 2 2023.
- [7] Ernst & Young. Environmental, health, safety and sustainability data and analytics Benefits to performance and the bottom line. Technical report, Ernst & Young, 2017.
- [8] Irving L Janis and Raymond H Fadner. A coefficient of imbalance for content analysis. *Psychometrika*, 8(2):1943, 1965.
- [9] Ram Ramanan. *Introduction to Sustainability Analytics*. CRC Press, 7 2018.
- [10] Nicolas Webersinke, Mathias Kraus, Julia Anna Bingler, and Markus Leippold. Climatebert: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010*, 2021.

## A Tutorial on Voxel-Based approaches for 3D Irregular Packing Problems

Alessio La Greca

KU Leuven, Dept. of Computer Science, NUMA, Ghent Campus

e-mail: [alessio.lagreca@kuleuven.be](mailto:alessio.lagreca@kuleuven.be)

Tony Wauters

KU Leuven, Dept. of Computer Science, NUMA, Ghent Campus

e-mail: [tony.wauters@kuleuven.be](mailto:tony.wauters@kuleuven.be)

3D Irregular Packing Problems (3DIPPs) are notoriously difficult optimization problems to solve. Regardless of the specific variant considered, the primary goal in such problems remains the same: to arrange a set of 3D items, which are not necessarily convex in shape and which can feature holes, into a 3D container. Applications for 3DIPPs can be found in many modern industries, such as Selective Laser Sintering, Diamond Cutting and Pallet Loading. One of the first challenges encountered by researchers addressing these problems is representational in nature and involves deciding how the geometric aspect of the problem should be confronted. Many representation methods for 3D items require a deep understanding of trigonometry, mathematics or other complex subjects, thereby presenting a steep barrier to entry for many researchers. How one represents 3D objects is a crucial decision as it influences how the two main constraints associated with 3DIPPs are handled: (i) the containment constraint, meaning all the selected items should be placed completely within the boundaries of the container, and (ii) the non-overlapping constraint, which involves ensuring that no two items within the container collide with one another. Moreover, if the representation method selected is unintuitive and challenging to use, it could significantly hinder progress with respect to the optimization algorithms built atop it. Using voxels to represent 3D irregular items appears to be a good trade-off between accessibility, robustness, and accuracy. Moreover, modern General Purpose GPUs can enhance the performance of voxel-based algorithms thanks to their intrinsic *Single Instruction Multiple Data* features. The barrier to entry for voxel-based implementations is low, but nevertheless requires a degree of intuition. Therefore, in this talk we will provide a tutorial on three different voxel-based approaches to confront the geometric challenge of 3DIPPs: one based on an explicit voxel-to-voxel collision test, a second based on Voxel Octrees, and a third that leverages the concept of Nofit Voxels. Our computational results demonstrate the positive impact of GPU-aided implementations, which are capable of providing significant speedups for computationally-intensive geometric queries. A short overview of a fourth approach to voxel-based collision detection, grounded in the use of the Discrete Fast Fourier Transform, will also be given.

# Improved upper bounds for the cage problem via exhaustive generation of lifts

L. Stubbe

KU Leuven campus Ghent, Department of Computer Science

e-mail: [louis.stubbe@kuleuven.be](mailto:louis.stubbe@kuleuven.be)

J. Jooken

KU Leuven KULAK, Department of Computer Science

e-mail: [jorik.jooken@kuleuven.be](mailto:jorik.jooken@kuleuven.be)

J. Goedgebeur

KU Leuven KULAK, Department of Computer Science

e-mail: [jan.goedgebeur@kuleuven.be](mailto:jan.goedgebeur@kuleuven.be)

T. Van den Eede

KU Leuven KULAK, Department of Computer Science

e-mail: [tibo.vandeneede@kuleuven.be](mailto:tibo.vandeneede@kuleuven.be)

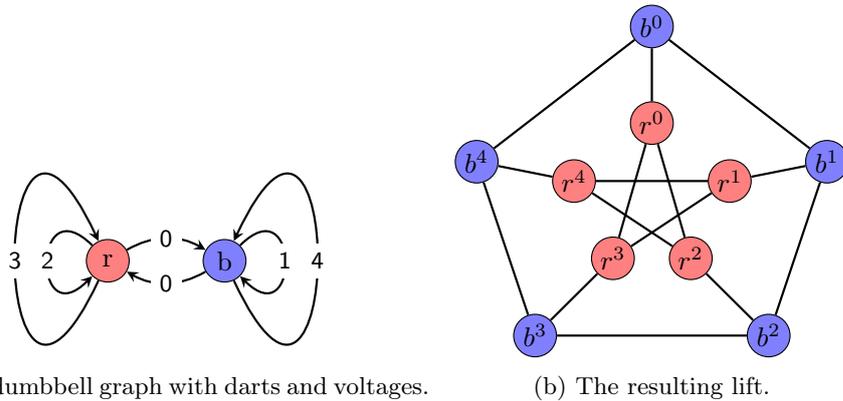
G. Exoo

Indiana State University, Department of Mathematics and Computer Science

e-mail: [ge@cs.indstate.edu](mailto:ge@cs.indstate.edu)

A  $(k, g)$ -graph is a  $k$ -regular graph with girth  $g$ , meaning the length of the shortest cycle is  $g$ . While the existence of such graphs for all  $k \geq 2$  and  $g \geq 3$  was established by Sachs [2], determining the minimum order  $n(k, g)$  remains a notoriously difficult challenge known as the *Cage Problem*. A  $(k, g)$ -graph of minimum order is called a  $(k, g)$ -cage. Figure 1b illustrates the only  $(3, 5)$ -cage, known as the Petersen graph. For most pairs  $(k, g)$ , the exact value of  $n(k, g)$  is unknown, and research focuses on narrowing the gap between the theoretical lower bounds and the best known upper bounds.

A common strategy to establish upper bounds is to constructively generate candidate graphs. Since many known small  $(k, g)$ -graphs exhibit a high degree of symmetry, we focus on generating highly symmetrical graphs directly using *lifts*. A lift requires three ingredients: a base graph  $G$  where every directed edge (dart) has an inverse, a finite group  $\Gamma$ , and a voltage assignment  $\alpha$  mapping darts to  $\Gamma$  such that  $\alpha(d^{-1}) = \alpha(d)^{-1}$ . The resulting lift, denoted  $G^\alpha$ , has vertex set  $V(G) \times \Gamma$ . The edges are defined such that a dart  $d$  from  $u$  to  $v$  in the base graph with voltage  $\alpha(d)$  induces a directed edge from  $u^s$  to  $v^{s \cdot \alpha(d)}$  for every  $s \in \Gamma$ . Figure 1 demonstrates this process, where the Petersen graph is constructed as a lift of the dumbbell graph using the cyclic group  $C_5$ .



(a) The dumbbell graph with darts and voltages.

(b) The resulting lift.

Figure 1: A lift of the dumbbell graph with the group  $C_5$  resulting in the Petersen graph. Opposing directed edges in the lift were replaced with a single undirected edge.

While lifts have been successfully utilized in the past to find new upper bounds for  $n(k, g)$  [1], these approaches were usually not exhaustive. We developed an algorithm to exhaustively generate lifts by exploring pairs of base graphs and groups, combined with a backtracking algorithm to determine voltage assignments. To make this computationally feasible, we implement aggressive pruning strategies based on the edge automorphisms of the base graph, automorphisms of the used group, and walks in the base graph. Using this method, we established new upper bounds for five cases of the Cage Problem. Additionally, we found six further improvements using excision techniques, yielding a total of eleven new upper bounds.

## References

- [1] G. Exoo and R. Jajcay. On the girth of voltage graph lifts. *European J. Combin.*, 32(4):554–562, 2011.
- [2] H. Sachs. Regular graphs with given girth and restricted circuits. *J. Lond. Math. Soc.*, s1-38(1):423–429, 1963.

## Exact Methods for the Vehicle Routing Problem with Time Windows and Stochastic Customers

Yulin Han

KU Leuven, ORSTAT

e-mail: [yulin.han@kuleuven.be](mailto:yulin.han@kuleuven.be)

Hande Yaman

KU Leuven, ORSTAT

e-mail: [hande.yaman@kuleuven.be](mailto:hande.yaman@kuleuven.be)

We study an extension of the vehicle routing problem with time windows by incorporating stochastic customers, i.e., ad-hoc service requests. The uncertainty in stochastic customers is formulated through scenarios. The problem is modeled based on a newly proposed a priori optimization framework and formulated as a scenario-based two-stage stochastic integer program. In the first stage, service routes are planned for regular customers, ensuring each is visited exactly once within their specified time windows. In the second stage, realized stochastic customers are either inserted between regular customers along the planned routes or outsourced at a predetermined cost through recourse actions. The objective is to minimize the total cost, comprising the routing cost for regular customers and the expected recourse cost associated with serving or outsourcing stochastic customers. To solve this problem, we propose two equivalent deterministic mixed-integer linear programming formulations and strengthen both formulations with valid inequalities. We develop two exact solution methods: a Cut-and-Branch (C&B) algorithm and an Integer L-shaped (ILS) algorithm. The ILS algorithm incorporates adapted optimality cuts based on the problem structure and several acceleration techniques. Computational results demonstrate the effectiveness of incorporating valid inequalities in tightening the linear relaxations of formulations and improving overall computational performance. The ILS algorithm shows superior scalability compared to the C&B algorithm as the number of scenarios increases. Furthermore, we evaluate the value of stochastic modeling and analyze the influence of insertion path flexibility on solution quality.

# A Markovian Mean-Field Model for the Dynamics of Societal Trust

Sindhu Chaudhary, Dieter Fiems, Koen De Turck

Ghent University, Department of Telecommunication and Information Processing

e-mail: {sindhu.chaudhary,Dieter.Fiems,Koen.DeTurck}@ugent.be

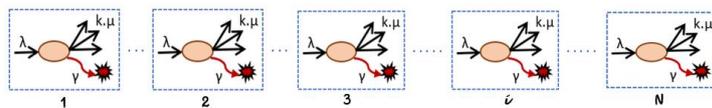
Trust is a core concept in the social sciences, structuring relations between government and citizens. It becomes particularly fragile during times of crises, when individual confidence is shaped not only by personal cognitive evaluation but also by social contagion and population-wide sentiment. To capture these intertwined mechanisms, we introduce a stochastic interacting-particle model for the evolution of trust within a large society.

At the individual level, trust is modeled as a birth-death process with catastrophic resets equivalent to an  $M/M/\infty$  queue with breakdowns. Arrivals and departures represent trust gains and losses, respectively, while a breakdown represents an event where a person completely loses trust. The arrival rate  $\lambda$  represents slow trust formation, while the departure rate  $n\mu$  (where  $n$  is the current trust level) captures loss-gain asymmetry motivated by prospect theory: losses weigh more heavily than gains, and higher trust levels carry greater potential for loss of confidence. The key feature is that the breakdown rate is not independent; rather, an individual is more susceptible to losing all trust when others have already done so. We capture this interdependence by letting the breakdown rate depend on the proportion of the population that has lost all trust ( $\pi_0$ ) in the government through the function  $\gamma(\pi_0)$ .

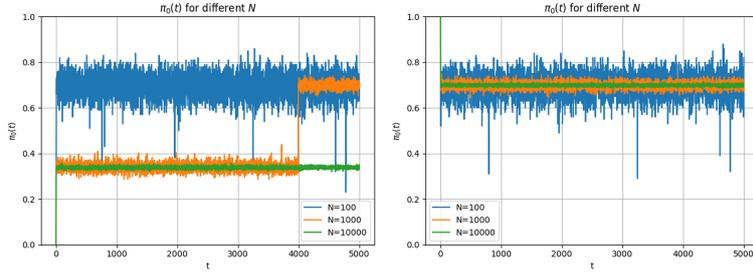
To study the interacting population model, we first solve the  $M/M/\infty$ -dynamics with breakdowns at the individual level (Figure 1). For constant breakdown rate  $\gamma$ , the stationary distribution  $(\pi_k)_{k \geq 0}$  admits a closed-form expression in terms of Kummer's confluent hypergeometric function  $M(a, b, z)$ ; in particular,  $\pi_0 = e^{-\lambda/\mu} M\left(\frac{\gamma}{\mu}, \frac{\gamma}{\mu} + 1, \frac{\lambda}{\mu}\right)$ . These single-agent results serve as building blocks for the interacting population model. To model social influence, we let the breakdown rate depend on the fraction  $\pi_0$  of individuals whose trust is entirely depleted. We focus on the logistic response, which is a flexible functional form widely used in behavioral and contagion modeling,

$$\gamma(\pi_0) = \gamma_0 + \frac{\gamma_1}{1 + \exp(-c(\pi_0 - \phi))}.$$

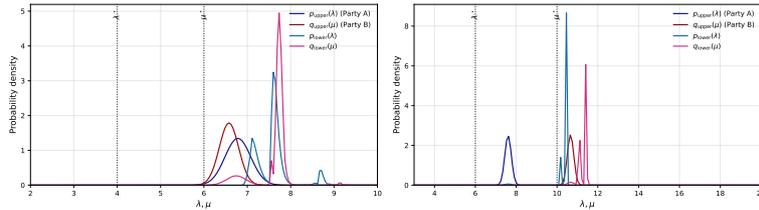
This breakdown rate ranges between  $\gamma_0$  and  $\gamma_0 + \gamma_1$  and increases with the share of fully distrustful individuals. The corresponding mean-field equations determine the time evolution and stationary values of the occupancy probabilities



**Figure 1:** Multiple individuals change states in parallel.



**Figure 2:** Observed Metastability in long trust regimes for some parameter values



**Figure 3:** Continuous Mixed Strategy Densities

$(\pi_k(t))_{k \geq 0}$  where  $k$  is the level of trust.

$$\frac{d\pi_0(t)}{dt} = \gamma(\pi_0(t))(1 - \pi_0(t)) + \mu\pi_1(t) - \lambda\pi_0(t),$$

$$\frac{d\pi_k(t)}{dt} = \lambda\pi_{k-1}(t) + (k+1)\mu\pi_{k+1}(t) - (\lambda + k\mu + \gamma(\pi_0(t)))\pi_k(t), \quad k \geq 1.$$

A notable feature emerging from this logistic-function-shaped feedback is metastability. For suitable parameter ranges, the system possesses two locally stable fixed points separated by an unstable saddle. Finite but large population systems may settle in either a relatively trusting configuration or a low-trust configuration. Stochastic simulations confirm that societal trust can switch between these regimes, with such switches becoming rarer as population size increases — the probability of transition decays exponentially with population size  $N$  (Figure 2). The mean switching time between metastable regimes can be approximated using WKB/large deviation theory.

Building on these results, we embed the trust dynamics into a game-theoretic framework. Model parameters are redefined as cost-controlled levers for a party in power and an opposition, who influence repair ( $\lambda$ ) and erosion of trust ( $\mu$ ). Using the Logit-Quantal Response Equilibrium (Logit-QRE), we compute equilibrium probability densities over strategies for both parties (Figure 3). This is done for both stable trust regimes whenever multiple fixed points for  $\pi_0$  are observed, yielding upper and lower  $\pi_0$  branches. The resulting densities illustrate how public trust can evolve and how strategic behavior interacts with breakdown feedback, offering insights for policymakers seeking to manage societal trust and prevent large-scale loss of confidence.

### Key references

- [1] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970.
- [2] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.
- [3] M. Assaf and B. Meerson. WKB theory of large deviations in stochastic populations. *Journal of Physics A: Mathematical and Theoretical*, 50(26), 2017.

# The Challenge of Quality Assessment in Dynamic Microsimulation

Morgane Dumont

University of Liège, HEC Liège Management School,  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [morgane.dumont@uliege.be](mailto:morgane.dumont@uliege.be)

Ahmed Alsaloum, Julian Ernst, Jan Weymeirsch, Ralf Münnich  
Trier University, Economic and Social Statistics Department

Spatial dynamic microsimulations probabilistically project geographically referenced units with individual characteristics over time [1]. Like any stochastic projection method, their outcomes are inherently uncertain and sensitive to multiple factors. In discrete time dynamic microsimulations, for each simulated time step (often years), each unit passes through different modules (births, deaths, ageing, partnership, employment, ...), evaluating if an event occurs for them via a Monte Carlo experiment. This inherently induces uncertainty due to the randomness of the method. However, simulations may also be sensitive to other factors, such as the choice of stochastic model types and complexity, as well as the parameter estimations among others.

Few articles detail the uncertainty in microsimulations, and the importance of its components is often overlooked. This is due to the high computational effort required for testing a large number of different simulation configurations and individual runs are necessary for the analysis. A complete sensitivity analysis testing the sensitivity to each single parameter in every module of the simulation would be unfeasible due to the complex structure of these microsimulations and the resulting computational power required to run them. Moreover, since microsimulations are typically developed to address specific problems and vary significantly in design and complexity, one-size-fits-all solutions are unattainable. Lastly, there is no commonly agreed-upon standard for reporting uncertainty in microsimulations [2].

Applying variance-based sensitivity analysis to both direct and indirect effects within the employment module of the MikroSim model for Germany [3], we show that commonly considered sources of uncertainty, namely coefficient and parameter uncertainty, are less influential than qualitative modelling choices. Because dynamic microsimulations are inherently complex and computationally intensive, it is crucial to consider potential factors of uncertainty and their influence on simulation outputs in order to more carefully design simulation setups and better communicate results. We find, that simple summary measures insufficiently capture overall model uncertainty and urge modellers to account for these broader sources when designing microsimulations and their results.

## References

- [1] Lovelace Robin, Dumont Morgane, *Spatial microsimulation with R*, CRC Press, Chapman & Hall/CRC The R Series, 2016.
- [2] Miia Rissanen, Jyrki Savolainen, *Uncertainty Analysis in Population-Based Dynamic Microsimulation Models: A Review of Literature*, Proceedings of the 16th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management - Volume 3: KMIS, 2024.
- [3] Münnich Ralf, Schnell Rainer, Brenzel Hanna, Dieckmann Hanna, Dräger Sebastian, Emmenegger Jana, Höcker Philip, Kopp Johannes, Merkle Hari-olf, Neufang Kristina, Obersneider Monika, Reinhold Julian, Schaller Jannik, Schmaus Simon, Stein Petra, *A population based regional dynamic microsimulation of Germany: the MikroSim model, methods, data, analyses - volume 15*, 2021

## A two-class queueing model for transplant waiting lists with renegeing high-priority patients

Thomas Debevere

Ghent University, TELIN

e-mail: [thomas.debevere@ugent.be](mailto:thomas.debevere@ugent.be)

Sabine Wittevrongel

Ghent University, TELIN

e-mail: [sabine.wittevrongel@ugent.be](mailto:sabine.wittevrongel@ugent.be)

Arnaud Devos

Ghent University, TELIN

e-mail: [arnaud.devos@ugent.be](mailto:arnaud.devos@ugent.be)

Joris Walraevens

Ghent University, TELIN

e-mail: [joris.walraevens@ugent.be](mailto:joris.walraevens@ugent.be)

### Abstract

We model a transplant waiting list as a discrete-time priority queue from which patients may renege. Two types of patients, namely high-priority and low-priority patients, arrive in the system. Organs arrive according to a Bernoulli process and are given to high-priority patients first; low-priority patients can only receive an organ in the absence of high-priority patients. High-priority patients each have a geometrically distributed deadline, upon expiration of which they succumb to their illness and renege from the system. By means of probability generating functions, we derive performance measures such as the mean numbers of waiting patients and mean waiting times for both priority classes, as well as the loss probability for high-priority patients.

**Keywords:** queueing theory, discrete-time, priority scheduling, renegeing.

We consider a discrete-time queueing model with two types of patients, namely class-1 (high-priority) patients and class-2 (low-priority) patients. In each slot, the numbers of incoming patients of both classes are generated as a random pair from a fixed joint distribution. Organ arrivals follow a Bernoulli process, with each time slot independently yielding an organ with probability  $\mu$ . When an organ arrives, it is given to the longest-waiting class-1 patient, or to the longest-waiting class-2 patient if no class-1 patients are present. An organ arriving to an empty system is assumed to be discarded.

While discrete-time priority queues have been studied extensively, notably by Walraevens [1], none of these previous models, to the best of our knowledge, account for renegeing. We model renegeing through the notion of “deadlines”, representing the maximum duration of time a patient stays in the system before renegeing if not served. To be precise, we define a patient’s deadline to be the number of slots counted from the one immediately after their arrival up to and including the slot in which they will renege if service does not occur first. Each high-priority patient samples their deadline, independently from the other patients, from the same (shifted) geometric distribution characterized by the parameter  $\alpha$ , referred to as the survival rate. This is mathematically equivalent to

each high-priority patient independently reneging with the same probability  $1 - \alpha$  in every slot. This reneging process was adopted from the work of Bruneel and Maertens [2], who modeled a single-class discrete-time queue with such a reneging process. In fact, since class-1 patients are unaffected by class-2 patients, the class-1 subsystem reduces to the model studied by Bruneel and Maertens when  $\mu = 1$ . The choice to model reneging for high-priority patients is motivated by the observation that, in practice, certain patients are often prioritized precisely because they have a higher probability of death if not treated in time.

We find a recursive functional equation for the joint probability generating function (pgf) of the numbers of waiting class-1 and class-2 patients in steady state. By solving this recursion, we derive separate expressions for the steady-state pgfs of the number of waiting class-1 patients and the number of waiting class-2 patients, which include infinite sums and products. While the mean number of waiting class-1 patients follows directly from the first derivative of the functional equation, obtaining the mean number of waiting class-2 patients requires a second differentiation, resulting in a substantially more intricate expression that we consider to be one of the main contributions of this work. Applying Little's law yields the mean waiting times for both priority classes, with the mean waiting time of class-1 patients unsurprisingly upper bounded by their mean deadline. We also provide an explicit expression for the loss probability of class-1 patients.

Numerical results demonstrate the stabilizing effect that reneging has on the system by reducing the class-1 load, increasing service opportunities for both the remaining class-1 patients and class-2 patients. This effect is particularly significant when service capacity is low, as illustrated in Figure 1.

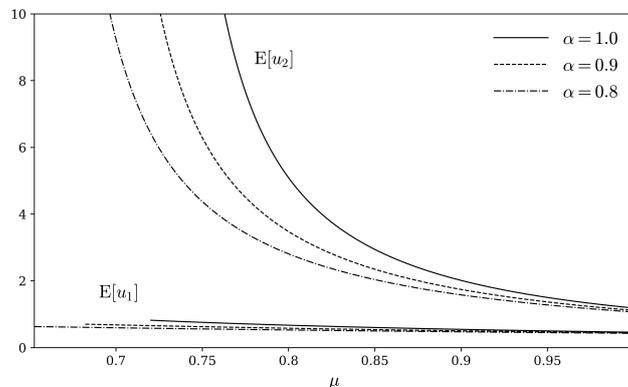


Figure 1: Mean numbers of waiting class-1 and class-2 patients versus the service rate  $\mu$  for various values of the class-1 survival rate  $\alpha$ .

## References

- [1] J. Walraevens, *Discrete-time queueing models with priorities*. PhD thesis, Ghent University, 2004.
- [2] H. Bruneel and T. Maertens, A discrete-time queue with customers with geometric deadlines. *Performance Evaluation*, 85-86:52-70, 2015.

# First Results on the Impact of Synchronizing Driver Breaks and Charging in Electric Long-Haul Routing

Omar El Habti

Hasselt University, Logistics

e-mail: [omar.elhabti@uhasselt.be](mailto:omar.elhabti@uhasselt.be)

Kris Braekers

Hasselt University, Logistics

e-mail: [kris.braekers@uhasselt.be](mailto:kris.braekers@uhasselt.be)

Ruben D'Haen

Hasselt University, Logistics

e-mail: [ruben.dhaen@uhasselt.be](mailto:ruben.dhaen@uhasselt.be)

Climate change poses increasing pressure on the transport sector, which currently accounts for a significant share of global greenhouse gas emissions. To support the transition to sustainable transportation, companies are exploring greener transport solutions, including electric vehicles (EVs). However, adopting EVs introduces new challenges to vehicle routing problems (VRPs), primarily due to limited driving ranges, long charging times and inadequate charging infrastructure. These constraints may hinder the competitiveness of EVs compared to internal combustion engine vehicles (ICVs), particularly in mid- and long-haul contexts. Additionally, long-distance freight operations are strictly governed by Hours-of-Service (HoS) regulations, enforcing mandatory rest and break periods to ensure road safety. These regulations, despite adding complexity, also create an opportunity for EVs. More particularly, mandatory driver breaks can be synchronized with charging operations, potentially enhancing the competitiveness of EVs.

This study investigates this opportunity. In last-mile logistics, modern EVs are more likely to have enough battery to finish routes without on-road charging. Since these routes are generally shorter, the likelihood of drivers needing to take mandatory HoS breaks also decreases. For those reasons the analysis in this study focuses on mid- and long-haul contexts, with average route distances of 500 km, where charging operations become more likely and total driving duration requires mandatory HoS breaks.

To analyze the impact of synchronization between driver breaks and EV charging, the study first models charging and break scheduling for a predefined route with a fixed customer sequence. For this setting an Electric Truck Driver Scheduling Problem (ETDSP) is introduced, and a mathematical model is presented that determines charging and break feasibility. The analysis then compares ICV and EV operations without HoS regulations, followed by scenarios that include mandatory breaks, to assess their influence on EV competitiveness. In a second step, the analysis was extended to a VRP context, which moves beyond fixed sequences by allowing multiple feasible route options and varying node combina-

tions. All experiments are conducted on a single-day planning horizon.

The findings provide insights into optimizing EV logistics and quantify how accounting for regulatory requirements in operational planning can reduce the cost gap between EVs and ICVs.

# Controlling Energy Consumption In Production Systems by Utilizing the Digital Twin

Siamak Khayyati

University of Liège, HEC-Liège Management School of the University of Liège

e-mail: [s.khayyati@uliege.be](mailto:s.khayyati@uliege.be)

Oktay Karabağ

Erasmus University Rotterdam, Erasmus School of Economics

e-mail: [karabag@ese.eur.nl](mailto:karabag@ese.eur.nl)

Berk Berkman

Koç University, College of Engineering

e-mail: [bberkman20@ku.edu.tr](mailto:bberkman20@ku.edu.tr)

Şebnem Demirtaş

Koç University, College of Engineering

e-mail: [sdemirtas20@ku.edu.tr](mailto:sdemirtas20@ku.edu.tr)

Barış Tan

Özyeğin University, Faculty of Engineering & Faculty of Business

e-mail: [baris.tan@ozyegin.edu.tr](mailto:baris.tan@ozyegin.edu.tr)

Given the environmental impact of industrial energy consumption, there is increasing pressure on manufacturers to decrease the energy consumption in their manufacturing plants. Recent studies have focused on this issue by considering controlling the energy mode of energy intensive resources in manufacturing. Namely, a lower energy consumption can be achieved by putting the machines in idle mode or turning them off when there is few items to process.

Controlling the energy mode of machines can be achieved optimally through threshold-type policies [1, 2]. These policies do not use the actual or predicted arrival times of items to the machine but use only aggregate data. The optimality of threshold-based policies partly depends on having memory-less inter-arrival time distributions. However, this assumption does not hold in all manufacturing settings. Moreover, digital twins create an opportunity for controlling the system with more information. Namely, the digital twin can use the current location of the items in the production system to simulate the future behavior of the system and give predictions for arrival times.

This work considers the problem of an automotive manufacturer that aims to decrease the energy consumption in its paint oven using a digital twin implementation. The paint oven is responsible for a considerable portion of energy consumption in the plant.

The location of the vehicles in the upstream can be used for generating arrival time scenarios and these scenarios can be used for scheduling the oven operations. Figure 1 depicts this process. As the digital twin is not able to predict the arrival times to the oven exactly, we develop a scenario-based mathematical programming approach to solve the problem.

Given the complexity of automotive manufacturing processes, many scenarios are typically needed to reflect the stochasticity in the system. However, solving the problem with many scenarios is not possible with commercial solvers currently. To remedy this, we establish multiple properties for the optimal solution to derive valid inequalities and propose a heuristic algorithm based on the genetic algorithm to solve the problem.

Our numerical results indicate that incorporating the digital twin is very effective in improving the performance of the system by allowing for decreasing the energy consumption by up to 30%. Additionally, our numerical results indicate that the valid inequalities are very effective for speeding up the process of finding the optimal solution. Our numerical results also show that for large problems, the heuristic method is able to identify close-to-optimal solutions very fast.

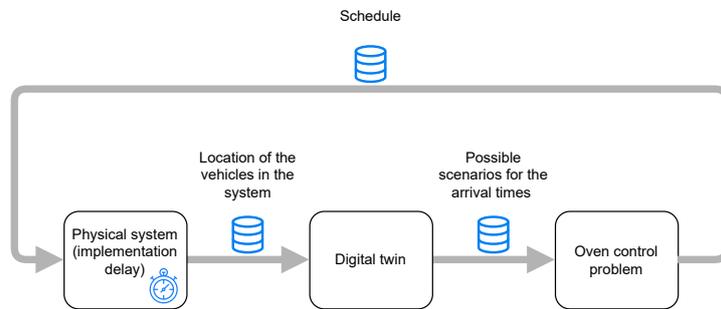


Figure 1: The flow of information between the physical system, digital-twin, and the oven control scheduler

## References

- [1] Tan, Barış and Karabağ, Oktay and Khayyati, Siamak, “Energy-efficient production control of a make-to-stock system with buffer-and time-based policies”, *International Journal of Production Research*, 2023.
- [2] Tan, Barış and Karabağ, Oktay, “A deterministic fluid model for production and energy mode control of a single machine”, *International Journal of Production Economics*, 2024.

# Assessing the Economic and Environmental Impacts of Peer-to-Peer Crowdshipping in Relation to Spatial Structure

Emma Innocente

UCLouvain, Center for Operations Research and Econometrics (CORE)

e-mail: [emma.innocente@uclouvain.be](mailto:emma.innocente@uclouvain.be)

Jean-Sébastien Tancrez

UCLouvain, Center for Operations Research and Econometrics (CORE)

e-mail: [js.tancrez@uclouvain.be](mailto:js.tancrez@uclouvain.be)

Freight transport accounts for approximately 8% of global greenhouse gas emissions (Greene, 2023), while it is projected to triple by 2050 compared to 2019 levels (Olivari et al., 2025). The rapid expansion of e-commerce has further intensified sustainability concerns, as the rise of direct-to-consumer deliveries tends to reduce freight consolidation, generating more frequent trips with smaller loads (Calabrò et al., 2023). In parallel, public and private institutions set clear targets for reducing greenhouse gas emissions (Olivari et al., 2025), prompting interest in alternative delivery models.

Among these alternatives, crowdshipping – where private individuals deliver parcels using personal vehicles along planned trips – has gained attention as a potentially cost-efficient and sustainable solution (Mohri et al., 2023; Pourrahmani & Jaller, 2021). In this system, crowdshippers accept to detour from their original route, within certain bounds, to pick up and deliver parcels in exchange for financial compensation (Archetti et al., 2016; Behrend & Meisel, 2018). By exploiting the unused capacity of existing personal mobility, crowdshipping aims to lower the need for dedicated delivery vehicles, thereby reducing delivery costs and emissions (Pourrahmani & Jaller, 2021).

However, knowledge on the actual impacts of crowdshipping remains limited (Buldeo Rai et al., 2018). The term itself covers a wide spectrum of models that differ by transport mode (e.g., car, bicycle, public transport), delivery configuration (e.g., direct, multi-leg, locker-based), and platform business models (Pourrahmani & Jaller, 2021; Mohri et al., 2023). Existing studies report mixed results, showing that the economic and environmental outcomes of crowdshipping are sensitive to several factors, among which we highlight: the spatial structure, the trip deviation, the matching mechanism, and the transport mode.

Building on these observations, this work evaluates the potential economic and environmental impacts of peer-to-peer car-based crowdshipping and identifies the spatial and operational conditions under which such impacts are observed. Particular attention is given to factors such as parcel demand, the number of participating crowdshippers, the sensitivity of mobility/logistics flows to distance,

and the spatial characteristics of the system – namely the hierarchy of city sizes and their spatial configuration.

Unlike much of the existing literature, we explicitly model the assignment of crowdshippers to parcels and account for the dynamic and uncertain nature of crowdshipping. To address this, we adopt an approximate dynamic programming (ADP) approach based on value function approximation (VFA), developed in previous work. This approach enables assignment decisions over time under uncertainty, accounting for stochastically arriving parcels and crowdshippers. We conduct a set of computational experiments, varying key parameters. This exploration helps identify the conditions affecting the economic and environmental performance of peer-to-peer car-based crowdshipping.

## References

- Archetti, C., Savelsbergh, M., & Speranza, M. G. (2016). The vehicle routing problem with occasional drivers. *European Journal of Operational Research*, *254*(2), 472–480. <https://doi.org/10.1016/j.ejor.2016.03.049>
- Behrend, M., & Meisel, F. (2018). The integration of item-sharing and crowdshipping: Can collaborative consumption be pushed by delivering through the crowd? *Transportation Research Part B: Methodological*, *111*, 227–243. <https://doi.org/10.1016/j.trb.2018.02.017>
- Buldeo Rai, H., Verlinde, S., & Macharis, C. (2018). Shipping outside the box: Environmental impact and stakeholder analysis of a crowd logistics platform in Belgium. *Journal of Cleaner Production*, *202*, 806–816. <https://doi.org/10.1016/j.jclepro.2018.08.210>
- Calabrò, G., Le Pira, M., Giuffrida, N., Fazio, M., Inturri, G., & Ignaccolo, M. (2023). A spatial agent-based model of e-commerce last-mile logistics towards a delivery-oriented development. *Transportation Research Interdisciplinary Perspectives*, *21*, 100895. <https://doi.org/10.1016/j.trip.2023.100895>
- Greene, S. (2023). Freight transportation. *MIT Climate Portal*. <https://climate.mit.edu/explainers/freight-transportation>
- Mohri, S. S., Ghaderi, H., Nassir, N., & Thompson, R. G. (2023). Crowdshipping for sustainable urban logistics: A systematic review of the literature. *Transportation Research Part E: Logistics and Transportation Review*, *178*, 103289. <https://doi.org/10.1016/j.tre.2023.103289>
- Olivari, E., Caballini, C., & Lluch, X. (2025). How to calculate GHG emissions in freight transport? A review of the main existing online tools. *Case Studies on Transport Policy*, *19*, 101343. <https://doi.org/10.1016/j.cstp.2024.101343>
- Pourrahmani, E., & Jaller, M. (2021). Crowdshipping in last mile deliveries: Operational challenges and research opportunities. *Socio-Economic Planning Sciences*, *78*, 101063. <https://doi.org/10.1016/j.seps.2021.101063>

# New developments in Artelys Knitro 15.1

Maxime Dufour

Artelys, [maxime.dufour@artelys.com](mailto:maxime.dufour@artelys.com)

Florian Fontan

Artelys, [florian.fontan@artelys.com](mailto:florian.fontan@artelys.com)

## 1 Introduction

Artelys Knitro is a mathematical programming solver for nonlinear and mixed-integer nonlinear problems. As input, it accepts linear structures, quadratic structures and black-box functions, with if possible, their first and second-order derivatives. Artelys Knitro relies on derivative-based algorithms to find locally optimal solutions. Knitro finds the global optimum for convex problems. For non-convex problems, Knitro converges to a first order stationary point (e.g. local optimum) for continuous models and is a heuristic for mixed-integer problems.

## 2 Recent improvements in Artelys Knitro 15.1

For nonlinear continuous problems, Artelys Knitro includes two interior point algorithms described in [1], a sequential linear quadratic programming algorithm and a sequential quadratic programming (SQP) algorithm. For mixed-integer nonlinear problems, Artelys Knitro includes a nonlinear branch-and-bound algorithm and a mixed-integer sequential quadratic programming algorithm.

Since Artelys Knitro 15 presented in [3], a new algorithm based on augmented lagrangian was added for continuous models. This algorithm is particularly suitable for complex problems with degenerate constraints where the linear independence constraint qualification (LICQ) is not satisfied. In this talk, we will present the characteristics of this new algorithm, how it complements the algorithms already available in Artelys Knitro and the applications on which it is particularly efficient. We will show the improvements on our benchmarks composed of both academic and industrial nonlinear instances.

## References

- [1] R. H. Byrd, J. Nocedal, and R.A. Waltz (2006). KNITRO: An integrated package for nonlinear optimization, In G. di Pillo and M. Roma, editors, *Large-Scale Nonlinear Optimization*, pages 35–59. Springer.
- [2] Ma, D., Saunders, M. A., & Orban, D. (2021). A Julia implementation of algorithm NCL for constrained optimization. In M. Al-Baali, L. Grandinetti, & A. Purnama (Eds.), *Numerical analysis and optimization: Vol. v* (pp. p–q). Springer International Publishing.

- [3] <https://www.artelys.com/news/knitro-15-0-new-tools-for-your-large-scale-models/>

## Comparing comparative process mining methods: towards a research agenda

Maxim Riebus

UHasselt - Hasselt University, Digital Future Lab

e-mail: [maxim.riebus@uhasselt.be](mailto:maxim.riebus@uhasselt.be)

Niels Martin

UHasselt - Hasselt University, Digital Future Lab

e-mail: [niels.martin@uhasselt.be](mailto:niels.martin@uhasselt.be)

Mieke Jans

UHasselt - Hasselt University, Digital Future Lab

e-mail: [mieke.jans@uhasselt.be](mailto:mieke.jans@uhasselt.be)

Comparing processes is a fundamental analytical task for organizations seeking to understand variation, performance differences, and opportunities for improvement. In domains such as healthcare, such comparisons are particularly relevant, for example when contrasting care pathways of different patient groups, examining how process execution differs over time periods, or benchmarking similar processes across hospitals or care networks. Process mining offers a data-driven foundation to support these types of comparisons by systematically analyzing event data that capture how processes are actually executed. Within the broader field of process mining, comparative process mining has been recognized as a distinct type of analysis, aimed at identifying differences and commonalities between process executions across variants, cohorts, or contexts [1].

Recent studies have highlighted several shortcomings in current comparative process mining methods. Firstly, existing techniques often narrowly focus on syntactic or control-flow aspects, neglecting perspectives such as probabilistic behavior, timing, or resource utilization [7]. Secondly, the definitions of what constitutes a relevant difference remain largely ad hoc, with little theoretical grounding. Finally, visualization techniques frequently fail to accommodate the complexity and multi-dimensionality of comparative findings [6].

The lack of standardization and methodological maturity in the field affects both theoretical progress and applicability. While methods such as log delta analysis enable verbalizing behavioral differences in natural language [3], they remain limited in scope and offer minimal support for multi-perspective or large-scale comparative settings. Likewise, recent methods that incorporate stochastic or statistical perspectives show potential, but they typically focus on localized or aggregate behavioral statistics, offering limited integration with an explicit process model that captures the global structure and relationships between process fragments [7].

To draw up a research agenda for comparative process mining, we propose

to review and compare existing comparative process mining methods. Over the past decade, a diverse set of methods has been developed to compare processes, often driven by specific analytical goals or application contexts. While these studies have demonstrated the potential of comparative process mining, they also reveal a fragmented methodological landscape, in which assumptions, comparison targets, and supported perspectives differ substantially across methods [2, 3, 5, 7, 6]. Limitations related to scope, interpretability, and perspective integration are frequently acknowledged in an ad hoc manner, but a systematic analysis that consolidates these insights across methods is still lacking. Given the growing relevance of comparative analysis within process mining, a structured comparison of existing methods is therefore needed to clarify their respective strengths and limitations and to create a research agenda for the further development of the field.

The need for improved support for comparative process mining is particularly evident in application domains where process comparison is a recurring analytical task, such as healthcare. In this context, process mining studies routinely compare care pathways across patient groups, time periods, or healthcare organizations, highlighting both the value of comparative insights and the practical difficulties of performing such analyses in a consistent and interpretable way with current methods [4]. By systematically analyzing existing comparative process mining methods, this work aims to contribute a clearer foundation that can better support process mining users and guide future advances in methods.

## References

- [1] van der Aalst WMP, Carmona J, editors. *Process Mining Handbook*. Lecture Notes in Computer Science, vol 448. Springer; 2022. doi:10.1007/978-3-031-08848-3.
- [2] Becker M, Laue R. A comparative survey of business process similarity measures. *Comput Ind.* 2012;63(2):148–167. doi:10.1016/j.compind.2011.11.003.
- [3] van Beest NRTP, Dumas M, García-Bañuelos L, La Rosa M. Log delta analysis: Interpretable differencing of business process event logs. *Lect Notes Comput Sci.* 2015;9253:386–405. doi:10.1007/978-3-319-23063-4\_26.
- [4] Partington A, Wynn M, Suriadi S, Ouyang C, Karnon J. Process mining for clinical processes: A comparative analysis of four Australian hospitals. *ACM Trans Manag Inf Syst.* 2015;5(4):1–18. doi:10.1145/2629446.
- [5] Hompes BFA, Buijs JCAM, van der Aalst WMP, Dixit PM, Buurman J. Detecting changes in process behavior using comparative case clustering. *Lect Notes Bus Inf Process.* 2017;244:54–75. doi:10.1007/978-3-319-53435-0\_3.
- [6] Wynn MT, Poppe E, Xu J, ter Hofstede AHM, Brown R, Pini A, van der Aalst WMP. ProcessProfiler3D: A visualisation framework for log-based process performance comparison. *Decis Support Syst.* 2017;100:93–108. doi:10.1016/j.dss.2017.04.004.
- [7] Mazhar TI, Tariq A, Leemans SJJ, Goel K, Wynn MT, Staib A. Stochastic-aware comparative process mining in healthcare. *Lect Notes Comput Sci.* 2023;14159:341–358. doi:10.1007/978-3-031-41620-0\_20.

## Federated Routing Optimization for the Physical Internet

Renaud De Landtsheer, Quentin Meurisse

Center of Excellence in ICT (CETIC - Charleroi)

Combinatorial Algorithmics (COAL) Department

{renaud.delandtsheer,quentin.meurisse}@cetic.be

Nowadays, despite the availability of optimization engines, 20% to 25% of the trucks moving in Europe are empty [1]. A typical example is a Belgian transporter that needs to transport a shipment (e.g., beers) to the south of France (e.g., because of warm weather). The truck can carry out the transportation, but it might come back empty. An obvious solution is to carry empty bottles back, but weight discrepancies lead to under-usage, and timing constraints apply. Another solution is to find a Transport Order (TO) for the return trip, typically via *Freight Marketplaces*, website listing TO's looking for transporter. Currently, these scenarios are common and essentially performed manually, in the sense that the transport companies must select the best TO's "by hand" through a computer database, a website, phone calls, mail exchange, etc.

To address these concerns, the concept of *Physical Internet* was theorized as a plan for a tight integration of logistics actors in order to improve the global efficiency of logistics and transportation systems. This notably involves a more global optimization of road transport.

### Federated Routing Optimization through Iterative Auction

We propose a technical approach to automatically consolidate truck routes and implement the scenario above using optimization tools without sharing private information.

The concept is formulated as an iterative auction where transporters can bid on TO's. At each iteration, there is a temporary assignment of TO's to transporters. Transporters bid on each TO they can see based on the marginal cost of serving this TO, knowing their temporarily assigned TO's. A TO is assigned to another transporter if his marginal cost for performing the TO is lower than the marginal cost of its currently assigned transporter. At each iteration, a transporter gives a single TO, takes a single TO, so that marginal costs are accurately taken into account. This process mimics a local search for VRP optimization. We can re-use concepts like ejection chains and the more efficient VLSN, allowing differential bids where a transporter declares that swapping one TO for another reduces their cost, and incorporate TO exchanges into the auction system.

The long term plan is to build a network of transporters and marketplaces where transporters can declare possible collaborations (e.g., to whom they can sub-contract transport orders) and can register to marketplaces. Transporters

therefore only see a fraction of all the TO's in the network and can only bid on the TO's they can see.

This system would actually be a federated routing optimization algorithm where the routing optimization engines operated by each transporter can optimize the routes of their trucks while trading TO's with each other through a standardized communication protocol.

The word *Federated* means that the optimization is performed without sharing sensitive data. Non-subcontracted TO's and truck routes are not disclosed to any central trusted third party node.

Gansterer notably developed an approach where TO's are first bundled based on their geography [2]. These bundles are put into single-shot auctions, i.e., they are definitively assigned to the cheapest transporter. This approach delivers significant speed and quality improvements compared to a single shot auction without bundling. However, it is only applicable if the auctioned TO's are published to all participating transporters. Furthermore, while this is adequate for last mile delivery, where there are many TO's to bundle together, it might not be sufficient for international transport across Europe due to the European *cabotage* rule, which constrains trucks to perform a maximum of six stops in a round trip (roughly speaking).

## Research Challenge and Requirements

The key technical research challenge is to develop the proper communication protocol and algorithms so that route optimization engines can connect with each other and perform route optimization while trading about transport orders. This protocol must address the following requirements:

**Efficiency, Optimality, Scalability:** The global algorithmic system must be designed so that the overall quality and the run time are acceptable at the scale of the market segment and accounting for communication delays. **Confidentiality:** As mentioned, the confidentiality of transporter's route and TO's is a strong requirement (GDPR, IA-act); the auction approach does not require disclosing this information directly. It is however possible that such a protocol indirectly leaks some information. **Pricing:** The result of the trading is that TO's are assigned to transporters. Transporters must be paid to perform them. The pay could either be the marginal cost declared in the algorithm, or the marginal cost by the second cheapest transporter as done in Vickrey auctions.

**Acknowledgement:** This research was financed by the service public de Wallonie Recherche under grant n°2010235 - ARIAC

## References

- [1] W. Ferrell, K. Ellis, P. Kaminsky, and C. Rainwater. *Horizontal collaboration: opportunities for improved logistics planning.*. International Journal of Production Research, 58(14):4267–4284, 2019.
- [2] M. Gansterer and R. F. Hart. *Request evaluation strategies for carriers in auction-based collaborations.* OR spectrum, 38:3–23, 2016

## Assessing the Value of Dual Sourcing for Production Continuity

Thomas De Backker

KU Leuven, Research Centre for Operations Management

e-mail: [thomas.debackker@kuleuven.be](mailto:thomas.debackker@kuleuven.be)

Joost van der Haar

KU Leuven, Research Centre for Operations Management

e-mail: [joost.vanderhaar@kuleuven.be](mailto:joost.vanderhaar@kuleuven.be)

Robert Boute

KU Leuven, Research Centre for Operations Management

e-mail: [robert.boute@kuleuven.be](mailto:robert.boute@kuleuven.be)

In the wake of recurring global disruptions, manufacturers increasingly invest in resilience measures to safeguard production continuity. Yet, the empirical, quantitative, component-level value of these interventions remains insufficiently examined. This study quantifies the operational and economic impact of dual sourcing implemented as backup sourcing, where a pre-established secondary supplier is activated when the primary supplier cannot immediately meet a firm's demand. Leveraging a large-scale dataset of stockouts for 1,009 components from a global original equipment manufacturer over the period between 2021 and 2024, the analysis examines whether, where, and when backup sourcing effectively reduces stockout risk and thereby enhances production continuity.

Building on the component stockout risk assessment by [1], stockout risk is decomposed into two operational drivers: the frequency of stockouts over a pre-defined time horizon and the average production downtime per stockout. The resulting expected number of stockout days serves as a proxy for production continuity and forms the primary outcome of interest to investigate the impact of dual sourcing. To isolate the causal effect of dual sourcing on these outcomes, the study adopts a doubly robust causal inference framework that combines outcome regression with inverse probability weighting, using XGBoost models. This approach creates a fair comparison between single- and dual-sourced components that share similar operational and design-related characteristics, such as lead time, design complexity, and consumption rates.

The empirical results show that dual sourcing leads to a statistically significant reduction in the expected number of stockout days per component, primarily through a shorter downtime once a stockout occurs, while the effect on stockout frequency is more modest. This pattern highlights that backup sourcing functions mainly as a recovery-focused resilience strategy: it does not prevent supply chain

disruptions, but helps in the production recovery once disruptions occur.

In addition, the effect of dual sourcing is heterogeneous, which can be interesting for targeting those components where it is most valuable. Components with long times, high consumption intensity, and complex designs or critical materials experience the largest improvements in production continuity when a backup supplier is available.

To bridge the gap between operational performance and financial decision-making, the study translates the estimated reduction in stockout days into a component-specific sourcing premium. This premium represents the economically justified willingness-to-pay for sourcing from a backup supplier and is expressed as the additional cost that can be afforded by the anticipated reduction in profit-at-risk resulting from fewer stockout days. The resulting cost–benefit framework enables procurement and operations managers to prioritize dual sourcing for components where the resilience payoff outweighs the additional sourcing cost.

Overall, the study contributes to the supply chain resilience literature by providing large-sample causal evidence on the value of a resilience intervention in a manufacturing environment. Methodologically, it illustrates how causal inference techniques can be used to move beyond intuition and stylized models toward data-driven assessments of resilience strategies at the component level. The findings offer guidance for component-level resilience decisions from both an operational and financial perspective, supporting more, evidence-based investments in supply chain resilience.

## References

- [1] T. De Backker, A. Vercammen, R. Boute, A data-driven component risk matrix to assess supply chain disruption risk, *International Journal of Production Economics* (2026) 109810.

# An Integrated Approach to Airport Security Rostering

Hicham Kamhi; Lotte Berghman

ORTEC B.V., Houtsingel 5 2719 EA Zoetermeer, The Netherlands

e-mail: [hicham.kamhi@ortec.com](mailto:hicham.kamhi@ortec.com)

## Abstract

This paper develops an integrated optimization framework for monthly rostering of airport ground-security personnel using a Simulated Annealing metaheuristic where a Minimum Cost Flow algorithm evaluates the solution's quality. The research addresses the complex problem of creating staff schedules while satisfying labor regulations, contractual constraints and employee preferences, using a practical dataset from real airport security operations.

## 1 Case study

In particular, airport security involves a set of predefined activities that must be continuously performed throughout the day. The security company receives staffing requirements for a single operational month exactly three months in advance, with requirements specified in half-hour time periods. An activity represents an operational task that needs staffing throughout the day, with personnel demand varying across different time periods. Each activity can only be performed by employees who possess the necessary skills.

Shifts are flexible, allowing them to start at any half-hour time period within the operational day and have a duration ranging from 4 to 9.5 hours in half-hour increments. This flexibility in both start time and duration is intended to allow closer alignment between staffing levels and fluctuating skill demands throughout the day.

The dataset covers a rostering period of 28 consecutive days. The workforce consists of 657 employees, of which 487 are internal. Employees have an average of 2,3 skills each, and the overall skill portfolio contains 7 unique skills. The total demand for these skills over the planning horizon amounts to 22510 working hours.

## 2 Methodology

This research addresses the complex problem of creating staff schedules while satisfying labor regulations, contractual constraints and employee preferences. We use a Simulated Annealing metaheuristic, where a Minimum Cost Flow algorithm evaluates the solution's quality.

The solution method simultaneously creates shifts and assigns them to employees rather than treating these two problems sequentially. Two construction heuristics generate initial solutions, which are improved through Simulated Annealing where two types of neighborhood structures are used. On the one hand, assignment searchers change who works which shifts by swapping  $x$ -day blocks between two employees, swapping two non-overlapping  $x$ -day blocks within one employee, or replacing an employee's  $x$ -day block with a single uniform shift type (including days off). Shift searchers, on the other hand, refine existing shifts by moving the break within the shift, shifting the entire shift one time period earlier or later, or extending/shortening the shift by one time period at either the start or the end. Staffing assessment is incorporated to efficiently match available employees to demand slots, ensuring that skill requirements are met throughout the planning horizon. By combining demand slots and employees with identical skills, one can model this Assignment Problem as the Minimum Cost Flow problem, and apply the Cost-Scaling Push-Relabel algorithm [1] which is in practice substantially faster than the Hungarian Method [2].

### 3 Results

The algorithm delivers high-quality solutions according to the defined key performance indicators. Results show excellent staffing coverage above 99%, with controlled over-staffing of less than 1%. Overtime is managed effectively, with about 3% of employees working overtime and average overtime per working employee under 0.5 hours, while taking into account employee preferences for work schedules and preferred working block lengths. Use of external staff is kept at an appropriate level around 6% to support internal capacity when necessary. Furthermore, we did a sensitivity analysis to compare rosters when one of the indicators like under-staffing, over-staffing, overtime, external staffing, or preferences, is valued higher. This shows that the framework is also interesting for other companies and datasets where prioritization is different.

In conclusion, the framework produces practical, high-quality rosters that meet coverage requirements while balancing operational needs with employee preferences. It offers a scalable and adaptable solution for airport security operations and can be applied to other industries requiring continuous staffing, providing a strong foundation for further improvement and real-world implementation.

### References

- [1] A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by successive approximation. Proceedings of the nineteenth annual ACM symposium on Theory of computing, 15(3):7–18, 1990.
- [2] H. W. Kuhn. The hungarian method for the assignment problem. Naval Research Logistics Quarterly, 2(1-2):83–97, 1955.

# Coordinated Scheduling of Equipment and Staff in Blood Component Production

Aleyna Gürsoy

University of Bergamo, Department of Management, Information  
and Production Engineering, Dalmine (Bg), Italy  
e-mail: [aleyna.gursoy@unibg.it](mailto:aleyna.gursoy@unibg.it)

Roberto Pinto

University of Bergamo, Department of Management, Information  
and Production Engineering, Dalmine (Bg), Italy  
e-mail: [roberto.pinto@unibg.it](mailto:roberto.pinto@unibg.it)

Pieter Smet

KU Leuven, Department of Computer Science, Gent, Belgium  
e-mail: [pieter.smet@kuleuven.be](mailto:pieter.smet@kuleuven.be)

Greet Vanden Berghe

KU Leuven, Department of Computer Science, Gent, Belgium  
e-mail: [greet.vanden.berghe@kuleuven.be](mailto:greet.vanden.berghe@kuleuven.be)

Ettore Lanzarone

University of Bergamo, Department of Management, Information  
and Production Engineering, Dalmine (Bg), Italy  
e-mail: [ettore.lanzarone@unibg.it](mailto:ettore.lanzarone@unibg.it)

The supply of blood and blood products is managed by the so-called Blood Supply Chain, which is composed of four echelons: collection, production, inventory, and distribution [1]. After donation, most whole blood units are separated into their components, mainly red blood cells, plasma and platelets. This separation ensures more efficient use of whole blood units collected from healthy donors.

This work focuses on the production echelon and proposes an advanced scheduler for the fractionation of blood products from whole blood units. Fractionation mainly consists of centrifuging the whole blood to separate its components based on density and physically separating them with the aid of an extractor, complemented by other supporting activities [2].

We apply a two-stage strategy, where the two stages correspond to the activities carried out before and after the so-called quarantine period, during which semi-processed blood units are refrigerated while awaiting the results of infectious disease tests. Each stage is modeled through a dedicated ILP formulation, and

the models are then solved applying an *ad hoc* matheuristic. The problems are then tackled using a rolling-horizon procedure, where delays in one stage reduce the resource availability in the subsequent one.

Previous research [3], which provided realistic and optimized makespans, assumed that an operator is partially engaged for the entire time the centrifugation or extraction process is in progress, while he/she actually intervenes only at the beginning and the end of these processes, to load and unload the blood bags, respectively. Partial engagement mitigated the impact of this assumption, but overestimated or underestimated operator utilizations were obtained.

In this work, the model formulation is restructured to assign tasks to operators in a more realistic way, engaging them for the parts of the processes in which they are actually involved.

The proposed methodology has been validated on data from the Immunohematology and Transfusion Medicine Department of the Niguarda Hospital in Milan, Italy, considering both the current production layout and possible future layouts.

Preliminary results confirm that the solutions are more realistic in terms of operator assignments and utilizations, but the computational effort required to solve the models increases dramatically. This will require modifying the models or developing different heuristics capable of scaling to real-world instances.

## References

- [1] Osorio, A.F., Brailsford, S.C., Smith, H.K., 2015. A structured review of quantitative models in the blood supply chain: a taxonomic framework for decision-making. *Int J Prod Res*; 53: 7191–7212.
- [2] Gürsoy, A., Pinto, R., Piccinini, F., Ghezzi, D., Veronese, L., Volpato, E., Rossini, S., Lanzarone, E., 2025. Production of blood components from donated units: literature review, areas for improvement, and research perspectives. *Flex Serv Manuf J*; DOI 10.1007/s10696-025-09645-2.
- [3] Gürsoy, A., Pinto, R., Piccinini, F., Ghezzi, D., Veronese, L., Volpato, E., Rossini, S., Lanzarone, E., 2026. A scheduling tool for the production of blood components. Paper under review.

## Optimizing Tactical Assembly Line Feeding in Picker-to-Part and Part-to-Picker Settings

Gohar Porbar

Ghent University, Department of Industrial Systems Engineering and Product Design  
FlandersMake@UGent — core lab ISyE, Ghent, Belgium  
e-mail: [gohar.porbar@ugent.be](mailto:gohar.porbar@ugent.be)

Emmanouil Thanos

Ghent University, Department of Business Informatics and Operations Management  
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium  
e-mail: [emmanouil.thanos@ugent.be](mailto:emmanouil.thanos@ugent.be)

Hannah Verplancke

Ghent University, Department of Business Informatics and Operations Management  
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium  
e-mail: [hannah.verplancke@ugent.be](mailto:hannah.verplancke@ugent.be)

Veronique Limère

Ghent University, Department of Business Informatics and Operations Management  
FlandersMake@UGent — core lab CVAMO, Ghent, Belgium  
e-mail: [veronique.limere@ugent.be](mailto:veronique.limere@ugent.be)

El-Houssaine Aghezzaf

Ghent University, Department of Industrial Systems Engineering and Product Design  
FlandersMake@UGent — core lab ISyE, Ghent, Belgium  
e-mail: [elhoussaine.aghezzaf@ugent.be](mailto:elhoussaine.aghezzaf@ugent.be)

In today's highly competitive and dynamic market, manufacturers face increasing pressure to produce a wide range of product variants while maintaining high service levels and cost efficiency. Mixed-model assembly lines have become a standard solution to meet this demand variability, as they allow different product variants to be assembled on the same line without major changeovers. However, this flexibility introduces new challenges in line feeding. In particular, the available space at the Border of Line (BoL) is limited, making it impractical to store large quantities of parts or a wide mix of components directly at the assembly line.

To cope with BoL space limitations, manufacturers rely on a variety of line feeding policies that differ in space usage, labor requirements and logistics effort. We consider five line feeding policies defined by Schmid and Limère [1]: Line Stocking (full pallets at the BoL), Boxed Supply (repacking parts into smaller boxes), Sequencing (presenting parts based on assembly order), Stationary Kit-

ting and Traveling Kitting (grouping parts for a single product). Assigning each part to an appropriate policy is a key tactical planning decision, as it determines the required preparation workload, the resulting material flow, and the consumption of BoL space. Line feeding decisions must also account for real-world constraints, including existing layout structures, limited preparation capacity, routing restrictions, and the available vehicle fleet. Efficiently managing these constraints is critical for finding a feasible solution for supplying parts to the line.

To capture different operational environments, we consider two settings. In the picker-to-part setting, preparation activities such as repackaging, sequencing, and kitting are carried out in a manual supermarket within dedicated cells with limited capacity. In the part-to-picker setting preparation is supported by an automated storage and retrieval system (AS/RS) that delivers the required part to logistic operators. The AS/RS imposes different operational constraints, such as the retrieval throughput, which influence the cost structure and feasibility of feeding policies, requiring a tailored optimization model for the automated setting. Our research proposes an exact solution approach determining the optimal assignment of feeding policies while minimizing logistics costs. We develop mixed-integer programming (MIP) models that explicitly incorporates the real-world limitations of the system.

We apply both models to a set of diverse instances to compare the operational performance of the two settings. The comparison highlights how the optimal mix of feeding policies changes with the preparation system. In the part-to-picker environment, reduced walking and search times lower the cost of labor-intensive policies such as kitting, making them more attractive. Overall, the results quantify cost savings from automation and identify conditions under which manual supermarkets remain competitive. This study provides manufacturers with a decision support tool for tactical line feeding planning by evaluating the impact of different feeding policies on logistics costs and space utilization. It can also provide practical insights regarding investments in automation, illustrating the potential benefits and trade-offs of transitioning from manual supermarkets to part-to-picker settings.

## References

- [1] Schmid, N. A. and Limère, V. (2019). A classification of tactical assembly line feeding problems. *International Journal of Production Research*, 57:7586–7609.

# A model-based heuristic for a three-dimensional loading vehicle routing problem with split pickups and time windows

Emeline Leloup, Célia Paquay and Thierry Pironet

University of Liège, HEC Liège Management School

QuantOM, Research Centre for Quantitative Methods and Operations Management

e-mails: [emeline.leloup@uliege.be](mailto:emeline.leloup@uliege.be), [cpaquay@uliege.be](mailto:cpaquay@uliege.be),

[thierry.pironet@uliege.be](mailto:thierry.pironet@uliege.be)

Keywords: Packing, Routing, Insert-and-Fix

The global reverse logistics market is expected to reach \$954.5 billion by 2029 [6], which highlights the importance of developing efficient logistics operations [3]. In particular, Logistics Service Providers (LSPs) are in charge of the collection and transportation of products, known as the first-mile collection problem, which is a specific part of reverse logistics. The LSP must determine a plan to collect and load boxes of their customers in an urban area and bring them back to a depot, while minimizing their total cost. Complex routing and loading challenges must be considered such as time windows, maximum working duration, box stability and reachability. In this work, we study the three-dimensional loading vehicle routing problem with split pickups and time windows (3L-SPVRP-TW), as presented in [1].

This work proposes a Mixed Integer Linear Programming (MILP) formulation to the 3L-SPVRP-TW, with constant travel durations. This formulation extends those existing in the literature by simultaneously taking into account additional real-life features such as split pickup, maximum working duration, and reachability constraints. Standard exact B&B optimization methods are expected to perform poorly over a restricted computational time.

We decided to adapt the insert-and-fix (I&F) matheuristic introduced by [2]. This matheuristic, able to use the potential offered by the MILP, consists in sequentially solving smaller MILPs. The first MILP is obtained by considering a subset of variables. In the second MILP, some variables of the first MILP are fixed to their values and a subset of variables is added. This process continues until all the variables are considered. I&F has been applied to lot-sizing and vehicle routing problems [2], the three-dimensional bin packing problem solely [4] and to a generalized multi-port container stowage planning problem [5]. To the best of our knowledge, I&F has never been applied to the problem integrating both routing and packing problems. We therefore extend the I&F to the 3L-SPVRP-TW, where the construction of the subproblems relies on the idea of inserting customers, step by step, inside the routes. We carry out some computational experiments to assess the impact of the I&F parameters and to determine the limitations of this method.

## References

- [1] Leloup, E., Paquay, C., Pironet, T., Oliveira, J.F., 2025. A three-phase algorithm for the three-dimensional loading vehicle routing problem with split pickups and time windows. *European Journal of Operational Research* 323, 1, 45–61
- [2] Liberalino, C.H.P., 2012. Problèmes de production avec transport des composants. Ph.D. thesis, Université Blaise Pascal- Clermont-Ferrand II.
- [3] Marriott, J., Bektaş, T., Leung, E.K.H., Lyons, A., 2025. The billion-pound question in fashion e-commerce: Investigating the anatomy of returns. *Transportation Research Part E: Logistics and Transportation Review* 194, 103904
- [4] Paquay, C., Limbourg, S., Schyns, M., Oliveira, J.F., 2018. MIP-based constructive heuristics for the three-dimensional bin packing problem with transportation constraints. *International Journal of Production Research* 56, 4, 1581–1592.
- [5] Parreño-Torres, C., Çalık, H., Alvarez-Valdes, R., Ruiz, R., 2021. Solving the generalized multi-port container stowage planning problem by a matheuristic algorithm. *Computers & Operations Research* 133, 105383.
- [6] Statista, 2022. <https://www.statista.com/statistics/1090465/reverse-logistics-market-size-worldwide/>. Retrieved on September 26, 2025

# Component effects of the Late Acceptance Hill Climbing metaheuristic

Maarten Van den Abeele, Kenneth Sörensen

ANT/OR – University of Antwerp Operations Research Group,

Department of Engineering Management,

Faculty of Business and Economics

e-mail: {maarten.vandenabeele,kenneth.sorensen}@uantwerpen.be

Late Acceptance Hill Climbing (LAHC) is a simple yet effective single-parameter metaheuristic for combinatorial optimisation. The method has achieved good results during competitions and it is also used in real-world projects. [1]

Despite empirical success, the components of LAHC and their interactions are understudied. The method differs from regular Hill Climbing (HC) as follows: it uses a history list, of which the length is the only algorithmic parameter. An acceptance mechanism then uses the historic values to decide on acceptance of candidate solutions 'later' on in the search.

This study makes modifications to the acceptance mechanism, the working of the history list, and the stopping criterion. Those modifications are based on suggestions and insights from the paper explaining the basic method [1], Step Counting [2], Custom Late Acceptance [3] and Diversified Late Acceptance [4]. First experiments are conducted on the same TSP instances as in [1], using fractional factorial design to produce varied implementations of the metaheuristic. Results are analysed using the Wilcoxon rank-sum test.

This ongoing study aims to clarify what are necessary principles in LAHC and how interacting components implement those principles to ensure acceptable run times and good quality of results. The algorithm below ensures convergence by having a maximum in the history list that is non-worsening. Choosing a list entry as is done on line 6 to then update the list on line 10 converges differently from randomly picking an entry for the same history length, preliminary experiments show. One possible outcome then is demonstrating whether one option is more favourable than the other with respect to available run time.

---

## Algorithm 1 General Purpose Late Acceptance Hill Climbing (Descent)

---

```

1: Construct solution  $s$ , define history length  $L$  and stopping criterion
2: Populate entire history list  $H$  of length  $L$  with  $f(s)$ 
3:  $i \leftarrow 0$ 
4: while stopping criterion not met do
5:   Generate candidate solution  $s'$ 
6:    $k \leftarrow i \bmod L$  ▷ Cycle through list
7:   if  $f(s') \leq H[k]$  then ▷ Original acceptance procedure
8:      $s \leftarrow s'$ 
9:   end if
10:   $H[k] \leftarrow f(s)$  ▷ Original list update
11:   $i \leftarrow i + 1$ 
12: end while

```

---

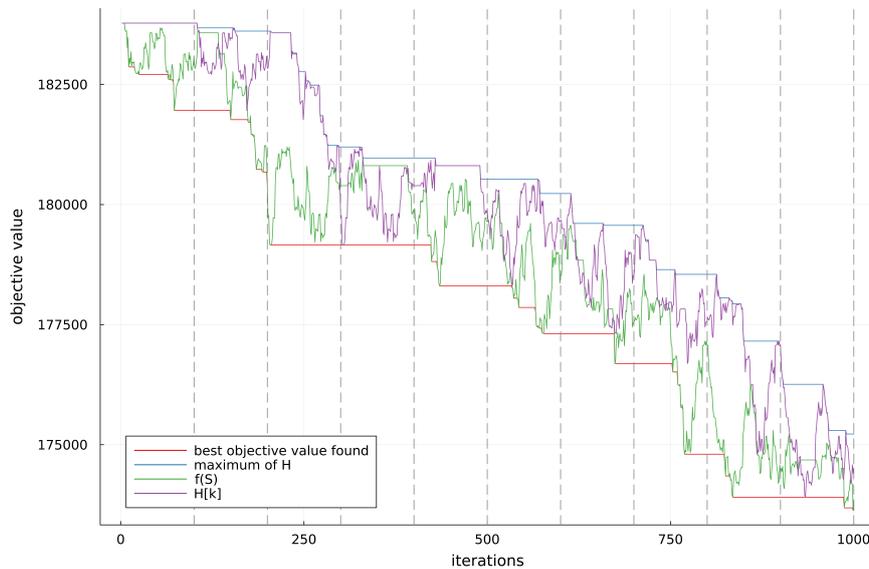


Figure 1: First 1000 iterations of Algorithm 1 applied to the rat783 instance ( $L=100$ , random construction, 2-opt)

## References

- [1] Edmund K. Burke and Yuri Bykov. The late acceptance hill-climbing heuristic. *European Journal of Operational Research*, 258(1):70–78, 2017.
- [2] Yuri Bykov and Sanja Petrovic. A Step Counting Hill Climbing Algorithm applied to University Examination Timetabling. *Journal of Scheduling*, 19(4):479–492, 2016.
- [3] Sylvain Clay, Lucien Mousin, Nadarajen Veerapen, and Laetitia Jourdan. CLAHC - custom late acceptance hill climbing first results on TSP. In *GECCO '21: Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1970–1973, Lille, France, July 2021. ACM.
- [4] Majid Namazi, Conrad Sanderson, M. A. Hakim Newton, Md Masbaul Alam Polash, and Abdul Sattar. Diversified late acceptance search. In Tanja Mitrovic, Bing Xue, and Xiaodong Li, editors, *AI 2018: Advances in Artificial Intelligence*, pages 299–311, Cham, 2018. Springer International Publishing.

# Emotion-driven customer satisfaction prediction in call center conversations

Thomas De Cleen

Vlerick Business School, Area Marketing

e-mail: [thomas.decleen@vlerick.com](mailto:thomas.decleen@vlerick.com)

Philippe Baecke

Vlerick Business School, Area Marketing

e-mail: [philippe.baecke@vlerick.com](mailto:philippe.baecke@vlerick.com)

Frank Goedertier

Vlerick Business School, Area Marketing

e-mail: [frank.goedertier@vlerick.com](mailto:frank.goedertier@vlerick.com)

15 December 2025

## 1 Introduction

Customer satisfaction is a crucial metric for evaluating call center performance, yet traditional survey methods suffer from low response rates and limited scalability, leading to blind spots in service quality management. Previous research that uses emotions to predict customer satisfaction in call center settings has several limitations. Most studies focus on a single type of emotion representation at either the turn or call level and rely on a single modality, typically text. Emotions are often operationalized in a coarse-grained manner (e.g., positive versus negative). In addition, prior work rarely examines the role of agent emotions or provides systematic comparisons across modeling choices, such as client-only versus client-and-agent emotions or unimodal versus multimodal features [1] [2] [3]. As a result, it remains unclear which emotion representations and modeling choices are most informative, limiting the practical deployment of emotion-driven satisfaction prediction models in real-world call center settings. To our knowledge, this is the first study to structurally compare multiple end-to-end methods for predicting customer satisfaction using multimodal turn-level and call-level emotion representations in a real-world call center setting. We investigate the trade-offs between turn-level and call-level emotion features, the predictive value of textual and audio emotion features and the contribution of call center agent emotion features.

Our framework is designed for seamless integration into existing contact-center systems: organizations that collect call transcripts, recordings, and post-call scores can train their own satisfaction predictors without any manual emotion annotation. By enabling reliable estimation for non-responders, this methodology promises to enhance service-quality monitoring and operational decision-making in real-world call centers.

## 2 Methodology

Our pipeline processes 25,008 conversations with self-reported customer satisfaction labels on a 0–10 scale (mean 8.34, SD 2.59,  $N = 25,008$ ) from a large utilities company. First, turn-level textual and acoustic emotion features are extracted from the transcriptions and audio signals using Transformer-based models. Using the MELD dataset [4], we finetuned a RoBERTa model [5] for text and a WavLM model [6] for audio emotion features. These features are then fused using a BiGRU-based multimodal emotion classifier. From the text, audio, and fusion models, we extract both the 7-class emotion predictions and the corresponding last hidden-layer embeddings. Next, different call-level feature sets such as weighted average emotions [1] and emotion dynamics [2] are calculated from the turn-level features. Lastly, we train a feedforward neural network to predict satisfaction using call-level features, while a bidirectional many-to-one LSTM is trained on the sequences of utterance-level emotion features.

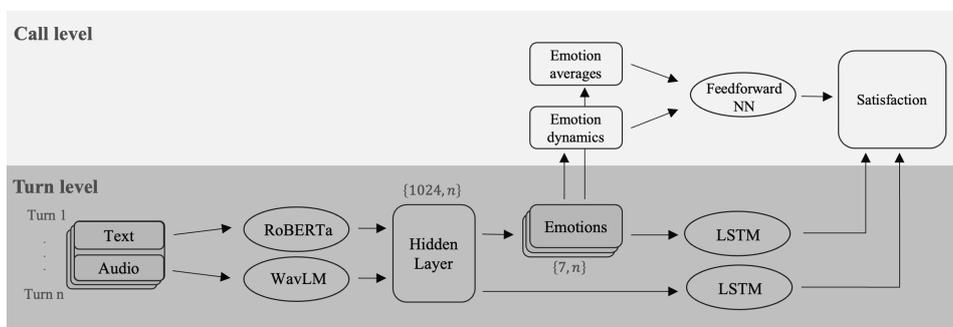


Figure 1: Methodology pipeline illustrating the use of turn-level and call-level emotion features for customer satisfaction prediction.

## References

- [1] De Cleen, T., Baecke, P., & Goedertier, F. (2025). The influence of emotions and communication style on customer satisfaction and recommendation in a call center context: An NLP-based analysis. *Journal of Business Research*, 189, 115192.
- [2] Guo, Y., Li, Y., Liu, D., & Xu, S. X. (2024). Measuring service quality based on customer emotion: An explainable AI approach. *Decision Support Systems*, 176, 114051.
- [3] Kim, Y., Levy, J., & Liu, Y. (2020). Speech Sentiment and Customer Satisfaction Estimation in Socialbot Conversations (No. arXiv:2008.12376). arXiv.
- [4] Poria, S., Hazarika, D., Majumder, N., Naik, G., Cambria, E., & Mihalcea, R. (2018). MELD: A Multimodal Multi-Party Dataset for Emotion Recognition in Conversations.
- [5] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pretraining Approach.
- [6] Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., ... & Wei, F. (2022). Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6), 1505-1518.

# A Bandit-Based Approach to Educational Recommender Systems: Contextual Thompson Sampling for Learner Skill Gain Optimization

Lukas De Kerpel

Faculty of Economics and Business Administration, Ghent University,  
FlandersMake@UGent, corelab CVAMO, Tweekerkenstraat 2, 9000 Ghent, Belgium  
e-mail: [lukas.dekerpel@ugent.be](mailto:lukas.dekerpel@ugent.be)

Dries F. Benoit

Faculty of Economics and Business Administration, Ghent University,  
FlandersMake@UGent, corelab CVAMO, Tweekerkenstraat 2, 9000 Ghent, Belgium  
e-mail: [dries.benoit@ugent.be](mailto:dries.benoit@ugent.be)

Arthur Thuy

Faculty of Economics and Business Administration, Ghent University,  
FlandersMake@UGent, corelab CVAMO, Tweekerkenstraat 2, 9000 Ghent, Belgium  
e-mail: [arthur.thuy@ugent.be](mailto:arthur.thuy@ugent.be)

January 29, 2026

In recent years, instructional practices have increasingly shifted toward Massive Open Online Courses (MOOCs) and other digital learning environments that provide online access to instructional materials, practice exercises, and assessments. Despite this digital shift, many courses continue to rely on standardized learning paths in which all learners follow the same sequence of exercises, offering limited opportunities for adaptation to individual learning needs. This challenge is particularly pronounced in MOOCs, where the scale of participation constrains instructors' ability to provide individualized guidance.

Educational recommender systems (ERS) offer a promising solution by supporting the **automated scaffolding** of exercises. Scaffolding is traditionally understood as an instructional practice in which a teacher provides structured support and gradually removes guidance as learners develop greater competence [1]. By dynamically guiding learners through sequences of exercises, ERS can align practice with learners' evolving knowledge states and provide adaptive support at scale.

This study introduces a contextual bandit framework for educational recommendation and presents the first empirical evaluation of Thompson Sampling (TS) in this setting. Specifically, we apply Linear Thompson Sampling (LinTS), a contextual bandit which models the expected learning benefit of exercises as a function of learner characteristics and past performance metrics. Within this

framework, recommendations are generated sequentially to select exercises that are most likely to advance a learner’s understanding of a targeted skill. Learning progress is quantified using a reward signal based on skill gain, defined as the change in a learner’s estimated knowledge state before and after completing an exercise. By optimizing directly for skill gain, the framework aligns recommendation decisions with sustained learning outcomes rather than commonly used short-term commonly used in ERS literature, such as correctness on the next exercise [2].

The framework is evaluated on the ASSISTments online tutoring dataset. Results show that LinTS consistently outperforms baseline approaches, yielding improvements in average skill gain of up to 20.7%.

In summary, the contributions of this work are threefold:

1. This study presents the first empirical evaluation of a TS-based contextual bandit within the domain of ERS.
2. The reward signal is based on skill gain, thereby aligning system optimization with sustained improvements in learners’ knowledge states rather than correctness on the next exercise.
3. The proposed framework demonstrates how contextual bandit models can be operationalized for large-scale educational data, providing a reproducible pipeline for sequential decision-making in online tutoring systems.

Overall, this work demonstrates how contextual bandit methods can enable effective, scalable, and interpretable scaffolding in digital learning environments.

## References

- [1] Janneke van de Pol, Monique Volman, and Jos Beishuizen, 2010, Scaffolding in Teacher–Student Interaction: A Decade of Research, *Educational Psychology Review*, 22(3), 271–296, <https://doi.org/10.1007/s10648-010-9127-6>.
- [2] Indu Manickam, Andrew S. Lan, and Richard G. Baraniuk, 2017, Contextual Multi-Armed Bandit Algorithms for Personalized Learning Action Selection, *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, <https://doi.org/10.1109/ICASSP.2017.7953377>.

## Reinforcement Learning for BESS short-term trading

Mouffe Mélodie

ENGIE Global Markets, Quantitative Research Modeling  
e-mail: [melodie.mouffe@engie.com](mailto:melodie.mouffe@engie.com)

Guillaume Brunieau

ENGIE Global Markets, Quantitative Research Modeling  
e-mail: [guillaume.brunieau@engie.com](mailto:guillaume.brunieau@engie.com)

Guillaume Fontaine

ENGIE Global Markets, Quantitative Research Modeling  
e-mail: [guillaume.fontaine1@engie.com](mailto:guillaume.fontaine1@engie.com)

Guillaume Keutgen

ENGIE Global Markets, Quantitative Research Modeling  
e-mail: [guillaume.keutgen@engie.com](mailto:guillaume.keutgen@engie.com)

Battery Energy Storage Systems (BESS) play an increasingly central role in enabling the flexibility required in modern electricity markets. Their economic performance, however, depends critically on the ability to make high-quality operational decisions under uncertainty, especially in highly volatile intraday environments.

Traditional optimization frameworks, typically linear or mixed-integer formulations, deliver fast and reliable schedules, but can be limited when confronted with rapidly evolving price signals, complex operational constraints, and the need for adaptive decision-making. Motivated by these challenges, the project presented here develops and evaluates a full reinforcement learning (RL) pipeline for short-term battery trading optimization.

A detailed intraday market simulation environment has been created, aligned with the organization's reference optimizer. It models battery states, operational costs, efficiency, imbalance penalties, market products, liquidity dynamics, and temporal constraints. This simulator enables controlled experimentation by providing forecasted and realized prices, battery trajectories, and market positions. The RL framework is implemented on top of this simulator, defining actions as the market bids, and ensuring rigorous benchmarking and consistency across training, validation, and backtesting.

The RL methodology is structured around a policy-based deep learning agent interacting with a multistep stochastic environment. The agent observes a com-

posite state including battery state-of-charge, traded energy, forecasted bid/ask signals, scheduled positions, and time-structured features.

Key modeling enhancements further strengthen the agent’s ability to make economically meaningful decisions, and new features have been introduced in the architecture to address the timeseries nature of price forecasts, which significantly improved stability and training efficiency compared with purely dense networks. A systematic program of hyperparameter analysis has been conducted to evaluate sensitivity to learning rates, price forecast horizons, decision granularity, entropy regularization, and training parallelization.

Finally, a hybrid decision-making approach that combines the strengths of reinforcement learning with the reliability of established optimization methods has been implemented. This design improves robustness and interpretability, while still benefiting from some uncertainty management.

The resulting RL system undergoes a comprehensive evaluation over several months of historical data. Backtesting includes rolling horizon decision-making, comparison against the existing optimization baseline, assessment of profit and losses, imbalance behavior, robustness to price volatility, and sensitivity analysis across multiple random seeds. Both training and testing phases reveal strong improvements over initial RL configurations, underscoring the importance of architecture choice. Several RL variants outperform the baseline optimization model on validation data and deliver competitive results on test periods.

Overall, this work demonstrates that carefully engineered reinforcement learning—combined with realistic simulation, robust modeling assumptions, and strong computational discipline—can serve as a credible decision-making tool for BESS optimization in volatile intraday environments. The study highlights both the promise and the practical challenges of deploying RL in industrial energy market applications, offering insights into algorithmic design, market modeling, and the balance between data-driven learning and operational reliability.

## Hexaly tutorial

Vianney Coppé, Lea Blaise

Hexaly

e-mail: [vcoppe@hexaly.com](mailto:vcoppe@hexaly.com); [lblaise@hexaly.com](mailto:lblaise@hexaly.com)

## Data-Driven Advanced-Planning-and-Scheduling Implemented in Printed-Circuit-Board Assembly

Chao Pan

KU Leuven, Research Centre for Operations Management,

e-mail: [chao.pan@kuleuven.be](mailto:chao.pan@kuleuven.be)

This study presents a Design Science Research initiative addressing the complex challenges of Printed Circuit Board assembly planning and scheduling in large-scale electronics manufacturing. Through a multi-year engagement with *TPV-Technology*, the world's leading monitor manufacturer, we developed and implemented a novel data-driven Advanced Planning and Scheduling system that integrates prediction, grouping and optimization. Field deployment across *TPV*'s 4 factories more than 3 years, demonstrated significant operational improvements: a 4% increase in overall equipment effectiveness, a 10% increase in throughput and an annual cost reduction exceeding 1 million euros. Our research makes practical and theoretical contributions. First, the successful deployment of our system across diverse factory contexts demonstrates its practical efficacy and establishes its transferability to other volatile and complex manufacturing environments. Second, our study contributes to DSR by showing how a data-driven APS system can be designed and sustained in dynamic manufacturing contexts. This is achieved through three mechanisms: (i) grouping bridges prediction and optimization; (ii) iteratively responding to evolving software and operational requirements; and (iii) extending the solution design to other factories or industries. These principles provide a replicable framework for developing solutions and bridge the gap between theoretical prediction-and-optimization and practical implementation.

## A DP–CP Hybrid Framework for Solving the Job-Shop Scheduling Problem

Emma Legrand

UCLouvain, ICTEAM

e-mail: [emma.legrand@uclouvain.be](mailto:emma.legrand@uclouvain.be)

Roger Kameugne

UCLouvain, ICTEAM

e-mail: [roger.kameugne@uclouvain.be](mailto:roger.kameugne@uclouvain.be)

Pierre Schaus

UCLouvain, ICTEAM

e-mail: [pierre.schaus@uclouvain.be](mailto:pierre.schaus@uclouvain.be)

Dynamic Programming (DP) [1] and Constraint Programming (CP) [9] are well-known frameworks used to solve the classical Job-Shop Scheduling Problem (JSP) [2]. In this work, we combine the capability of DP to decompose problems into small and overlapping subproblems to be solved recursively with the power of the CP filtering algorithms in the JSP. Based on the current state and the current decision, a CP model is constructed, and the transition state is obtained after propagating the constraints of this model.

A JSP is a classical optimization problem that consists of  $n$  jobs and  $m$  machines. Each job contains  $m$  operations and we denote  $o_{i,j}$  the  $j^{\text{th}}$  operation of job  $i$  where  $j \in \{1, \dots, m\}$  and  $i \in \{1, \dots, n\}$ . The set of operations is denoted  $I$  and each operation  $k \in \{1, \dots, n \times m\}$  (where  $k$  is the index of  $o_{i,j}$  in  $I$ , i.e.,  $k = i \times m + j$ ) has a duration  $p_k$  and a machine assigned  $m_k$ . A preorder is imposed on operations of the same job and two operations sharing the same machine cannot overlap. The goal of the JSP is to find the schedule of all operations of each machine, minimizing the total execution time.

The DP formulation of the JSP is adopted from [5]. A state is a tuple  $s = \langle est, done \rangle$  where

- $est$  is an  $n \times m$  array containing the earliest start time of each operation.
- $done$  is the set of operations already scheduled.

A decision variable  $d_k$  is the start time of operation  $k$  and is used to specify the decision taken by  $k$  in its domain  $D_k$ . In the transition state  $t_k(s_k, d_k)$ , the operation  $k$  is scheduled at its earliest starting time and set as a scheduled operation. The transition cost function  $h(s, d_k)$  is the maximum completion time of the scheduled operations.

In the transition state, the earliest starting time of operations is updated from the transit state. To do so, from the current state, the CP variables are initialized by setting the earliest starting time of unscheduled operations and the start time for scheduled operations. The precedence and resource constraints for each machine are propagated until the fixed point is reached, and the earliest starting time of unscheduled operations is updated. For resource constraints, the CP model

uses the global constraint *No-Overlap* described in [6]. This constraint encapsulates many filtering algorithms like edge-finding, not-first/not-last, detectable precedences, and overload checking. In [8, 7], it is proved that the edge-finding algorithm performs adjustments similar to immediate selection on an ascendant set introduced in [3, 4]. Additional adjustments of other filtering algorithms allow the combination to subsume the immediate selection on an ascendant set used in [5] for the DP model of JSP.

This presentation describes the proposed articulation of the DP-CP hybrid framework.

## References

- [1] Richard Bellman, 1954, The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6), 503–515.
- [2] Jacques Carlier and Eric Pinson, 1989, An algorithm for solving the job-shop problem. *Management Science*, 35(2), 164–176.
- [3] Jacques Carlier and Eric Pinson, 1990, A practical use of Jackson’s preemptive schedule for solving the job shop problem. *Annals of Operations Research*, 26(1), 269–287.
- [4] Jacques Carlier and Eric Pinson, 1994, Adjustment of heads and tails for the job-shop problem. *European Journal of Operational Research*, 78(2), 146–161.
- [5] Joaquim A. S. Gromicho, Jelke J. van Hoorn, Francisco Saldanha-da-Gama, and Gerrit T. Timmer, 2012, Solving the job-shop scheduling problem optimally by dynamic programming. *Computers & Operations Research*, 39(12), 2968–2977.
- [6] Petr Vilím, 2007, Global constraints in scheduling. *Univerzita Karlova, Matematicko-fyzikální fakulta*.
- [7] Paul Martin and David B. Shmoys, 1996, A new approach to computing optimal schedules for the job-shop scheduling problem. *Proceedings of IPCO, Lecture Notes in Computer Science*, vol. 1084, 389–403.
- [8] Philippe Baptiste and Claude Le Pape, 1995, A theoretical and experimental comparison of constraint propagation techniques for disjunctive scheduling. *Proceedings of IJCAI*, 600–606.
- [9] Francesca Rossi, Peter van Beek, and Toby Walsh, 2008, Constraint programming. *Foundations of Artificial Intelligence*, 3, 181–211.

## Collaborative multimodal freight transportation under uncertainty

François Gilet

UCLouvain, CORE

e-mail: francois.gilet@uclouvain.be

Jean-Sébastien Tancrez

UCLouvain, CORE

e-mail: js.tancrez@uclouvain.be

Stefan Creemers

UCLouvain, CORE

e-mail: stefan.creemers@uclouvain.be

December 4, 2025

Multimodal transportation has been identified as an important lever to reduce truck dominance in the freight transport sector (often referred to as “modal shift”) [5]. Reducing the share of goods transported by truck and using alternatives such as railway or inland waterways is expected to yield economic, environmental, and social benefits [3, 6], and is an objective of the EU [2]. However, multimodal transportation has seen little adoption in the past [4]. Tactical planning difficulties and long booking windows have been identified as important barriers [1, 4]. Our research aims at better understanding how horizontal collaboration, in the form of capacity pooling, can mitigate those barriers and increase the multimodal share by reducing the overall uncertainty.

In this work, we model a simple network with two hubs connected by road and rail, where companies facing demand uncertainty can book train capacity together. Once the uncertainty is resolved, collaborating companies must allocate the shared capacity between them. Each company continues to manage alone its pre- and post-haulage transport and must address any insufficient allocation of rail capacity by means of trucks. We assume the truck-only route is more expensive than the multimodal one.

On the one hand, if the total realized demand is below train capacity, a cost is unnecessarily incurred for each reserved unit that is not used. On the other hand, if the total realized demand exceeds train capacity, each unit above capacity must take the more expensive truck-only route. However, since the cost of the truck-only route and the pre- and post-haulage costs are different for each company, the allocation scheme of train capacity influences the total cost. We use a priority allocation rule where a company can only receive train capacity when all companies which would benefit more from multimodal transport are fully served. This scheme minimizes the total cost for the coalition, but will require the introduction of a profit sharing mechanism.

We develop a method to find the optimal joint train capacity booking for a single period. Inspired by the classic newsvendor problem, the total expected cost is minimized using the expected marginal value of each  $q$ -th reservable unit of train capacity, which is weighted by the probability distribution of each product and the allocation rule described above. Our method differs from the existing literature on the multi-product newsvendor problem by searching for a single shared order quantity.

Preliminary results show that collaboration can yield a significant reduction in costs and can increase the share of multimodal transport when compared to the setting in which companies do not collaborate. We explore the impact of various cost and demand parameters on the performance of our scheme.

## References

- [1] Shafagh Alaei, Koen Mommens, Javier Durán-Micco, and Cathy Macharis. Evaluating logistics companies' readiness towards adopting synchronomodality in the flanders region. *Sustainability (Switzerland)*, 16, 6 2024.
- [2] European Commission. White Paper on Transport: Roadmap to a Single European Transport Area - Towards a Competitive and Resource Efficient Transport System, 2011.
- [3] Pedro Furtado and Jean Marc Frayret. Proposal sustainability assessment of resource sharing in intermodal freight transport with agent-based simulation. In *IFAC-PapersOnLine*, volume 28, pages 436–441, 5 2015.
- [4] Ahmed Karam, Anders Julius Klejs Jensen, and Mohamed Hussein. Analysis of the barriers to multimodal freight transport and their mitigation strategies. *European Transport Research Review*, 15, 12 2023.
- [5] Ercan Kurtuluş and İsmail Bilge Çetin. Analysis of modal shift potential towards intermodal transportation in short-distance inland container transport. *Transport Policy*, 89:24–37, 4 2020.
- [6] Bram J. De Moor, Stefan Creemers, and Robert N. Boute. Breaking truck dominance in supply chains: Proactive freight consolidation and modal split transport. *International Journal of Production Economics*, 257, 3 2023.

## Balancing service levels and operational effort in dynamic smart waste collection

Dilay Aktas Dejaegere

KU Leuven, Institute for Mobility (LIM)

e-mail: [dilay.aktas@kuleuven.be](mailto:dilay.aktas@kuleuven.be)

Pieter Vansteenwegen

KU Leuven, Institute for Mobility (LIM)

e-mail: [pieter.vansteenwegen@kuleuven.be](mailto:pieter.vansteenwegen@kuleuven.be)

Cities face increasing pressure to keep waste collection both efficient and responsive to citizen needs. Although some waste streams, like recyclables, can generate revenue, the primary goal of municipal waste collection is to provide a service to the public. Traditionally, this service is carried out by fleets of heavy vehicles operating on fixed routes with predetermined collection frequencies. From a routing and optimization perspective, waste collection can be classified into arc-routing and node-routing problems. Arc-routing focuses on efficiently traversing streets, as in curbside collection, whereas node-routing targets servicing specific locations, such as public or underground bins. This work focuses on the latter.

The placement and sizing of the bins are typically determined through tactical planning: smaller bins are cheaper and space-efficient, while larger bins reduce collection frequency but require higher investment. The main challenge is that waste generation is inherently unpredictable, varying with many contextual factors such as population density, demographics, urbanization, environmental regulations as well as weather, holidays, school schedules, and events. As a result, fixed collection schemes can lead either to overflowing bins and forced clean-ups or to unnecessary collections in order to avoid overflows.

Smart bins equipped with sensors offer a promising way to address this gap [1]. Sensors can provide real-time or daily snapshots of fill levels and thus reduce uncertainty at selected locations. In this work, we present a dynamic optimization framework for smart waste collection aiming to balance service level and operational effort. Service level is measured by the number of realised daily bin overflows, reflecting the inconvenience to citizens. Operational effort is measured by vehicle-kilometres travelled, which affects driver workload, fuel consumption, environmental impact, and the traffic burden caused by heavy collection vehicles. While these two objectives inherently conflict, not all trade-offs are of practical interest. Municipalities might tolerate a small number of overflows, but collection plans that include very low vehicle-kilometres at the expense of widespread service failure would not be acceptable. For this reason, we adopt the perspective of maximising service level under a global budget of operational effort.

We develop a Mixed Integer Program (MIP) that determines which bins to collect on which day and the routes of the vehicles over a multi-period plan-

ning horizon. The model uses expected daily fill rates for all bins and incorporates daily sensor readings. The framework supports dynamic collection planning through daily re-optimisation as new information becomes available. However, we observe that under a global operational budget, frequent re-optimisation can disrupt the intended balance between service level and operational effort by repeatedly re-allocating resources. We therefore propose a conditional re-optimisation approach, where replanning is triggered by service level violations, in contrast to classical rolling-horizon methods.

We first evaluate the dynamic system under full sensor coverage to quantify the potential benefits of dynamic planning under operational budgets. This serves as a benchmark for service quality and operational efficiency. Since equipping every bin with a sensor is often not realistic [2], we also consider partial sensor coverage, where only a subset of bins provide sensor readings and the remainder rely on predicted daily fill rates. We analyse several baseline sensor placement strategies as in [3] and assess their performance relative to full sensor coverage. Because the system does not rely on probabilistic models but instead uses only expected fill rates and sensor readings for monitored bins, we believe it reflects the practical information constraints faced by cities and waste operators prior to large-scale sensor deployment.

The main contributions of this study are as follows. First, we develop a dynamic optimisation framework for smart waste collection that relies solely on expected fill rates and sensor readings, avoiding probabilistic assumptions to reflect the limited information typically available prior to large-scale sensor deployment. Second, we introduce a service-oriented planning perspective that maximises service level under a global operational budget, highlighting the challenges of resource allocation in a dynamic and uncertain environment. Third, we also analyse scenarios with partial sensor coverage and address the resulting need for proactive servicing of unmonitored bins, in order to avoid purely reactive collections.

## References

- [1] Lopes, M. and Ramos, T. R.P. *Efficient sensor placement and online scheduling of bin collection*. *Computers & Operations Research*, 151:106113, 2023.
- [2] Hess, C., Dragomir, A. G., Doerner, K. F., and Vigo, D. *Waste collection routing: a survey on problems and methods*. *Central European Journal of Operations Research*, 32(2):399–434, 2024.
- [3] Cubillos, M., Spliet, R. and Wøhlk, S. *On the effect of using sensors and dynamic forecasts in inventory-routing problems*. *INFOR: Information Systems and Operational Research*, 60(4):473–490

## How we scheduled ORBEL: new models, takeaways and feedback

Felix Rauh

KU Leuven

e-mail: [felix.rauh@kuleuven.be](mailto:felix.rauh@kuleuven.be)

In this talk, we share behind-the-scenes insights into how this conference was scheduled. Building on the participant-preference-based approach introduced by Vangerven et al. [1], we present the original model, our adaptations, key implementation considerations, and how the final schedule reflects these ideas. We also share some fun data insights and invite participants to provide feedback.

The approach by Vangerven et al. follows a three-phase structure: first deciding which talks run in parallel to minimize conflicts, then grouping them into blocks corresponding to sessions, and finally assigning sessions to timeslots while respecting presenter availability. This "minimize-conflicts-first" strategy optimizes globally but may separate talks that audiences want to see together. We therefore also explored an alternative "coherent-sessions-first" approach that begins by pairing talks with shared audience interest, ensuring popular combinations stay in the same session, before assembling them into parallel slots. This can be further extended by requiring paired talks to share keywords for added thematic coherence.

Beyond the algorithm, we analyzed the submitted preferences to uncover fun insights: which talk was most popular, and which "dream team" pair of talks was most frequently liked together. Did participants in the survey distribute their votes equally over all talks or did nobody scroll until the end? We present these preference-based awards and conclude with a live audience survey to evaluate the scheduling quality and understand how participants actually use the conference program. We hope this talk provides useful insights and plan to upload a ready-to-use implementation for future conference organizers most likely under [https://github.com/felixrauh/conference\\_scheduler](https://github.com/felixrauh/conference_scheduler).

## References

- [1] Vangerven, B., Ficker, A.M.C., Goossens, D.R., Passchyn, W., Spijksma, F.C.R., Woeginger, G.J., 2018. Conference scheduling A personalized approach. *Omega* 81, 3847. <https://doi.org/10.1016/j.omega.2017.09.007>

# Generative AI as a Copilot for Mathematical Optimization Modeling

Lennart Lahrs

Gurobi Optimization

e-mail: [lennart.lahrs@gurobi.com](mailto:lennart.lahrs@gurobi.com)

Generative AI (GenAI) tools are increasingly capable of supporting the formulation of mathematical optimization models. This presentation explores how GenAI can assist modelers in translating problem descriptions into formal models, refining constraints and objectives, and iterating on alternative formulations. Rather than replacing human expertise, GenAI is positioned as a modeling copilot that enhances productivity while preserving model transparency and rigor. Practical examples illustrate both the potential benefits and current limitations of GenAI-supported mathematical modeling in Operations Research.

# The Incomplete Traveling Tournament Problem (iTTP)

David Van Bulck

Ghent University, Department of Business Administration and Operations Management

e-mail: [david.vanbulck@ugent.be](mailto:david.vanbulck@ugent.be)

Karel Devriesere

Ghent University, Department of Business Administration and Operations Management

e-mail: [karel.devriesere@ugent.be](mailto:karel.devriesere@ugent.be)

Dries Goossens

Ghent University, Department of Business Administration and Operations Management

e-mail: [dries.goossens@ugent.be](mailto:dries.goossens@ugent.be)

In a double round-robin sports tournament, each team meets every other team once at home and once away. One of the most extensively studied scheduling problems in this setting is the Traveling Tournament Problem (TTP). It minimizes the total travel distance of all teams by grouping consecutive away games into a single road trip, allowing teams to travel directly between away venues rather than repeatedly returning home. Literature on the TTP is vast, ranging from the development of dedicated lower bounds (e.g., [4]) over algorithm design (e.g., [2]) and analysis of approximation algorithms (e.g., [5]).

As an alternative to the round robin format where each team meets every other team a fixed number of times, recently the incomplete round-robin tournament format where teams meet only a subset of the other teams has gained popularity. Perhaps most known is the renewed competition format of the UEFA Champions League group stage, which follows the incomplete format. But also in amateur leagues, the format is gaining in popularity (see e.g., [1] for an application to Belgian hockey).

This presentation introduces a new optimization problem which we coin the incomplete traveling tournament problem (iTTP), which as the name suggest introduces the classic TTP into the realm of incomplete tournaments. Given an even numbered set of teams and rounds and pairwise distances between the teams' home venues, the problem is to schedule an incomplete round robin tournament such that *(i)* each team plays precisely one game per round, *(ii)* each team plays an equal number of home and away games, *(iii)* no team plays more than three home or away games in a row, *(iv)* each team meets every other team *at most* once, and *(v)* the overall distance traveled by all teams is minimized.

Apart from introducing several problem instances for the iTTP, we show why a natural generalization of the independent lower bound for the classic TTP is weak for the iTTP and we show how to strengthen this bound. Moreover, we

show how a trivial IP formulation and constructive heuristic based on [3] fare on this problem.

## References

- [1] Devriesere, K. and Goossens, D. (2025). Redesigning Belgian youth field Hockey competitions using an incomplete round-robin tournament. *INFORMS Journal on Applied Analytics*, 55, 457-468.
- [2] Di Gaspero, L., and Schaerf, A. (2007). A composite-neighborhood tabu search approach to the traveling tournament problem. *Journal of Heuristics*, 13:189-207.
- [3] Li, M., Van Bulck, D., and Goossens, D. (2025). Beyond leagues: A single incomplete round robin tournament for multi-league sports timetabling. *European Journal of Operational Research*, 323:208-223.
- [4] Urrutia, S., Ribeiro, C. C. and Melo, R. A. (2007). A new lower bound to the traveling tournament problem. *IEEE Symposium on Computational Intelligence in Scheduling (SCIS'07)*, 15-18.
- [5] Westphal, S. and Noparlik, K. A 5.875-approximation for the traveling tournament problem (2012). *Annals of Operations Research*, 218:347-360.

## Novel neighborhood structures for incomplete round robin tournaments

Karel Devriesere

Ghent University, Department of Business Informatics and Operations Management

e-mail: `karel.devriesere@ugent.be`

David Van Bulck

Ghent University, Department of Business Informatics and Operations Management

e-mail: `david.vanbulck@ugent.be`

Dries Goossens

Ghent University, Department of Business Informatics and Operations Management

e-mail: `dries.goossens@ugent.be`

In this work, we study neighborhood structures for incomplete round robin tournaments. We are given a set of teams  $T$  with  $|T|$  even and a set of rounds  $R$  such that each team plays exactly one match in each round and faces the same opponent at most once. This setting naturally arises in contexts where  $|T|$  is much larger than  $|R|$ , for example in youth sports competitions [1, 2]. Moreover, the timetable should be balanced, meaning that a team should play at least  $\lfloor \frac{|R|}{2} \rfloor$  and at most  $\lceil \frac{|R|}{2} \rceil$  home games.

A popular way to obtain sports timetables is to use local search methods. Neighborhood structures, however, have so far solely been studied in the context of round robin tournaments, where  $|R| = \lambda(|T| - 1)$  with  $\lambda \geq 1$ , i.e., where each team sees every opponent exactly the same number of times. These neighborhood structures are known as *Round Swap*, *Partial Round Swap*, *Team Swap*, *Partial Team Swap* [3] and *Cycle reversal* [4]. In this work, we fill this gap by investigating whether and how these neighborhood structures can be applied to incomplete round robin tournaments. Moreover, we propose two novel neighborhood structures called *Incomplete Partial Team Swap* and *Incomplete Partial Round Swap*. In presenting these neighborhood structures, we make a connection with graph theory and show that Incomplete Partial Round Swap is connected for certain values of  $|R|$ , meaning that starting from any timetable, we can reach any other timetable by a finite sequence of neighborhood moves. Moreover, we embed the neighborhood structures in a Late Acceptance Hill Climbing heuristic and test the algorithm on three sets of instances from the literature. We show that the algorithm produces high quality timetables for all three instance sets.

## References

- [1] Devriesere, K. and Goossens, D. (2025). Redesigning Belgian youth field hockey competitions using an incomplete round-robin tournament. *INFORMS Journal on Applied Analytics*.
- [2] Li, M., Van Bulck, D., and Goossens, D. (2025). Beyond leagues: A single incomplete round robin tournament for multi-league sports timetabling. *European Journal of Operational Research*, 323(1):208–223.
- [3] Januario, T., Urrutia, S., Celso, C. R., and de Werra, D. (2016a). Edge coloring: A natural model for sports scheduling. *European Journal of Operational Research*, 254:1 – 8.
- [4] Knust, S. and von Thaden, M. (2006). Balanced home–away assignments. *Discrete Optimization*, 3(4):354–365

## Fair Policies for Sequential Resource Allocation

Sten Wessel

TU Eindhoven, Department of Mathematics and Computer Science

e-mail: [s.wessel@tue.nl](mailto:s.wessel@tue.nl)

Christopher Hojny

TU Eindhoven, Department of Mathematics and Computer Science

e-mail: [c.hojny@tue.nl](mailto:c.hojny@tue.nl)

Frits Spieksma

TU Eindhoven, Department of Mathematics and Computer Science

e-mail: [f.c.r.spieksma@tue.nl](mailto:f.c.r.spieksma@tue.nl)

We consider the following *sequential resource allocation problem*. There is a (single) divisible resource with an initial supply  $s$ , which has to be divided over  $n$  clients. Clients are visited sequentially, in a predetermined order, identified as  $[n] := \{1, \dots, n\}$ . Each client has a certain *demand*, which is the amount of the resource they want to receive from the supplier. However, the demands are unknown to the supplier beforehand; it is only revealed when arriving at the client. Instead, a probability distribution for the demand of every client is known to the supplier in advance. Let  $D_i$  for every client  $i \in [n]$  denote the random variable modeling the demand of client  $i$ . We assume that the demand distributions for every client are *independent, discrete*, and take a finite number of values. The initial supply  $s$ , the set of demand distributions  $\{D_i : i \in [n]\}$  for every client, together with a specific objective define the Sequential Resource Allocation (SRA) problem. A solution to the SRA problem comes in the form of a *policy* that determines the amount  $x_i$  of the resource that is allocated to client  $i \in [n]$  upon visiting that client, which may depend on the allocation history of clients earlier in the sequence.

SRA serves as a basic online allocation problem with stochastic demand. In practice, the problem arises with different objectives and additional properties depending on the context. For example, Bassok and Ernst [1] are one of the first to consider the SRA problem when maximizing revenue. However, in many non-profit applications like distributing food over food bank locations, the main objective is that an allocation should be *fair*. In the remainder of this article, we focus on maximizing fairness and consider the *fill rate* of an allocation as a measure of fairness. The fill rate  $\varphi_i$  of client  $i \in [n]$  is the ratio  $\varphi_i = x_i/d_i$  of the allocation  $x_i$  and their actual, realized demand  $d_i$ .

We consider two objectives related to fairness of allocations, which we refer to as *ex-post* and *ex-ante* fairness. For the ex-post objective, the aim is to maximize the expected minimum fill rate of all clients, i.e., maximize

$$\mathbb{E}_{(d_1, \dots, d_n) \sim (D_1, \dots, D_n)}[\min\{\varphi_i : i \in [n]\}]. \quad (1)$$

In contrast to maximizing the expected minimum fill rate, *ex-ante* fairness maximizes the minimum expected fill rate of all clients, i.e.,

$$\min\{\mathbb{E}_{(d_1, \dots, d_n) \sim (D_1, \dots, D_n)}[\varphi_i] : i \in [n]\}. \quad (2)$$

Several policies for SRA have been proposed. We highlight Lien et al. [2] aiming to achieve *equity* of fill rates among clients, motivated by supplying regional and local food banks in Chicago in the United States, and Manshadi et al. [3], who study the SRA problem motivated by allocating scarce medical supplies to hospitals during the COVID-19 pandemic.

A salient feature of all policies proposed for SRA is that the position of a client in the sequence may have a profound impact on the amount of resource allocated to the client. Indeed, when considering a setting where all clients are identical, i.e., have the same distribution  $D_i = D$ , results from [2] imply that for several policies the amount of resource allocated to clients depends heavily on the position in which the client is served. That is, although existing policies aim to achieve fairness by maximizing the minimum fill rate, individual clients could experience these policies as unfair if they receive less than another client that is identical. From a client’s perspective, it would thus be favorable to have an SRA policy that is also *individually fair*: the property that identical clients should be treated equally. The topic of individual fairness has not been considered before in the SRA literature.

In this work, we introduce a new allocation policy, based on linear programming, that is individually fair. Further, we provide an extensive computational study to evaluate existing policies and our new policy with respect to ex-post and ex-ante fairness. Using a fine-tuned implementation of existing methods, we are able to evaluate significantly larger problem instances—with over 20 clients—than those previously considered in the literature, which were limited to at most 6 clients. As a consequence, we are able to assess how the performance of various policies scales with the number of clients. We are also able to show to what extent the position of a client in the sequence affects their allocation.

We find that our new allocation policy performs worse on the ex-post fairness objective, whereas it outperforms existing methods on ex-ante fairness when the supply of the resource is scarce. Furthermore, the average fill rate is much more equal for all clients in the sequence compared to any of the methods, where effects of position in the sequence are clearly noticeable.

## References

- [1] Yehuda Bassok and Ricardo Ernst, 1995, Dynamic allocations for multi-product distribution. *Transportation Science*, 29(3):256–266.
- [2] Robert W. Lien, Seyed M. R. Iravani, and Karen R. Smilowitz, 2013, Sequential resource allocation for nonprofit operations. *Operations Research* 62(2):301–317. DOI: 10.1287/opre.2013.1244.396.
- [3] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz, 2023, Fair dynamic rationing. *Management Science*, 69(11):6818–6836, DOI: 10.1287/mnsc.2023.4700.

# Smart Lotteries in School Choice: Ex-ante Pareto Improvement with Ex-post Stability

Tom Demeulemeester

Maastricht University, Department of Quantitative Economics

e-mail: [tom.demeulemeester@maastrichtuniversity.be](mailto:tom.demeulemeester@maastrichtuniversity.be)

Haris Aziz

University of New South Wales, School of Computer Science and Engineering

e-mail: [haris.aziz@unsw.edu.au](mailto:haris.aziz@unsw.edu.au)

Péter Biró

Eötvös Loránd University (ELTE), Centre of Economic and Regional Studies

Corvinus University, Department of Operations Research and Actuarial Sciences

e-mail: [biro.peter@krtk.elte.hu](mailto:biro.peter@krtk.elte.hu)

Gergely Csáji

Eötvös Loránd University (ELTE), Centre of Economic and Regional Studies

Corvinus University, Department of Operations Research and Actuarial Sciences

e-mail: [csaji.gergely@krtk.elte.hu](mailto:csaji.gergely@krtk.elte.hu)

We study the school choice problem, where school seats are assigned to prospective students who have preferences over them. Because schools are often indifferent between large subsets of the students, it is common practice to use random tie-breaking techniques. In this paper, we study how we can improve the students' assigned schools, in comparison to applying the celebrated deferred acceptance (DA) algorithm with random tie-breaking, without violating the stability of the solution. We show that both this problem and computing the expected outcome of DA with random tie-breaking are NP-hard. We also propose a column generation procedure to minimize the expected rank over all lotteries over stable matchings, and illustrate how it improves upon DA with random tie-breaking, as well as upon other existing mechanisms from the literature that Pareto-dominate DA.

# Two-Step Imbalanced PU Learning with Graph Neural Networks

Bérard Romain

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [romain.berard@kuleuven.be](mailto:romain.berard@kuleuven.be)

Aerts Matthias

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [matthias.aerts@skynet.be](mailto:matthias.aerts@skynet.be)

Bruno Deprez

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

University of Antwerp - imec, Department of Mathematics

e-mail: [bruno.deprez@kuleuven.be](mailto:bruno.deprez@kuleuven.be)

Eduardo Ortega Vázquez Carlos

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [carlooseduardo.ortegavazquez@kuleuven.be](mailto:carlooseduardo.ortegavazquez@kuleuven.be)

Positive-Unlabelled (PU) learning addresses classification settings where only confirmed positives are observed, while the remaining examples are unlabelled and form an unknown mixture of positives and negatives. This setting is common in high-stakes applications such as risk screening and fraud detection, where the absence of a label does not imply a true negative (Bekker & Davis 2020). On relational data, the challenge is amplified by dependencies between samples: nodes are connected through interactions or transactions, and the graph structure itself carries predictive signal. Graph neural networks (GNNs) exploit this structure by learning node representations via neighbourhood aggregation, combining attributes with relational context (Hamilton et al. 2018).

PU learning is further complicated by the mechanism through which positives are labelled. Under *Selected Completely At Random* (SCAR), positives are labelled with a constant propensity, independent of their attributes. Under *Selected At Random* (SAR), the labelling propensity varies with observed information, so observed positives may be biased toward more detectable cases. This is often the more realistic regime in fraud detection: investigations tend to uncover the most detectable fraudsters, while less detectable cases remain hidden.

Existing graph-based PU approaches typically assume SCAR, balanced classes, and rely on full-batch objectives. These assumptions limit applicability to large, evolving, and highly imbalanced graphs, particularly when the observed positives are selected with a non-uniform propensity.

We propose GNN-NNIF, a scalable two-step framework for imbalanced PU node classification designed for transductive graphs and SAR mechanisms. In

Step 1, we sample random node subgraphs (Hamilton et al. 2018) to compute neighbourhood-aggregated representations and train a GNN encoder. Before any encoder training, we perform a first pass that only aggregates information to obtain initial embeddings. On these embeddings, we apply a Nearest-Neighbour Isolation Forest (NNIF) (Ortega Vázquez et al. 2023) to identify hidden positives in the unlabelled nodes. NNIF first computes anomaly scores with an Isolation Forest, and then adjusts these scores using a  $k$ -NN refinement based on distances to labelled positives. The highest-scoring unlabelled nodes are treated as hidden positives and removed from the provisional negative pool, yielding reliable negatives. We then train the encoder on labelled positives and reliable negatives with a joint objective: a label-propagation loss that encourages consistent posteriors in subgraphs, and a soft contrastive loss that increases separation between nodes with dissimilar inferred posteriors while placing emphasis on uncertain regions and the minority class. In Step 2, we re-apply NNIF to the learned embeddings to further refine the reliable-negative set. We then address imbalance by applying ADASYN oversampling and training an XGBoost classifier on the final node embeddings.

We evaluate GNN-NNIF on multiple citation benchmarks and a highly imbalanced fraudulent transaction graph under both SCAR and SAR labelling mechanisms, spanning a range of imbalance ratios. Preliminary results indicate that GNN-NNIF is consistently competitive with state-of-the-art PU baselines, with the strongest gains in highly imbalanced regimes.

## References

- Bekker, J. & Davis, J. (2020), ‘Learning from positive and unlabeled data: a survey’, *Machine Learning* **109**(4), 719–760.  
**URL:** <https://doi.org/10.1007/s10994-020-05877-5>
- Hamilton, W. L., Ying, R. & Leskovec, J. (2018), ‘Inductive Representation Learning on Large Graphs’. arXiv:1706.02216 [cs, stat].  
**URL:** <http://arxiv.org/abs/1706.02216>
- Ortega Vázquez, C., vanden Broucke, S. & De Weerd, J. (2023), ‘A two-step anomaly detection based method for PU classification in imbalanced data sets’, *Data Mining and Knowledge Discovery* **37**(3), 1301–1325. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 3 Publisher: Springer US.  
**URL:** <https://link.springer.com/article/10.1007/s10618-023-00925-9>

# ML-Based B2B Quotation Prioritization: Evidence from a Randomized Field Experiment

Bert Vandendriessche

Ghent University, Department

e-mail: `bervdndr.vandendriessche@ugent.be`

Dries Benoit

Ghent University, Department

e-mail: `dries.benoit@ugent.be`

Costly sales cycles and limited resources often force business-to-business (B2B) firms to prioritize open quotations. Typically, this happens in a subjective way. Machine Learning based prioritization has proven to outperform those subjective judgments when it comes to predicting win probability. No empirical evidence exists on the true added value of using such predictions as decision support in a real world environment. This study contributes to closing this gap by researching and empirically evaluating a machine learning-based quotation prioritization mechanism through a real-world field experiment. Our results reveal that data-driven prioritization significantly improves conversion rates and profitability, and a decrease of costly B2B sales cycle duration. These findings underscore the potential of machine learning as a decision-support tool in B2B lead management.

## On the Difficulty of Detecting Cherry-Picking in Counterfactual Explanations

James Hinns<sup>1</sup>      Sofie Goethals<sup>1</sup>      Stephan Van der Veecken<sup>1</sup>  
 Theodoros Evgeniou<sup>2</sup>      David Martens<sup>1</sup>

<sup>1</sup>University of Antwerp      <sup>2</sup>INSEAD

Artificial intelligence (AI) increasingly shapes decisions that affect individuals, organisations and society at large. Yet the complexity of modern AI systems often makes them difficult to interpret, which poses a substantial obstacle to adoption and effective oversight [2, 6]. In response, the field of Explainable Artificial Intelligence (XAI) has emerged, developing methods that make the decision-making processes of AI systems more transparent [1, 7].

Explanations serve several purposes: they can improve user trust, support accountability, and enable compliance with regulatory requirements. For example, the General Data Protection Regulation (GDPR) establishes a putative “right to explanation” for individuals affected by algorithmic decision-making, while the EU AI Act mandates explainability in high-stakes applications [4].

Despite rapid progress, explanation methods can yield substantially different results for the same prediction, even when using the same method. This phenomenon is known as the *disagreement problem* [3, 5]. The resulting plurality of valid explanations enables *cherry-picking*, where an explanation provider selectively presents only those explanations that best serve their interests.

Consider a fictional example. A model is used for loan decisions based on three features: income, gender, and years worked. Alice applies for a loan and is rejected. She then exercises her right under the GDPR to request an explanation. Bob, the explanation provider, generates several counterfactual explanations and presents the following: “If you had earned €10,000 more and worked two more years, you would have been accepted.” This appears reasonable to Alice, who later receives a promotion and is accepted upon reapplying. However, another valid counterfactual also exists: “If you had identified as male instead, you would have been accepted.” By presenting only the first explanation, Bob conceals the model’s biased behaviour. The withheld alternative reveals that gender influences the decision, while the selected explanation appears neutral and avoids exposing this issue. This selective presentation constitutes cherry-picking.

Our goal is twofold: (a) to formally define what constitutes a cherry-picked explanation, and (b) to characterise when cherry-picking can be detected, and what information is required to do so. To investigate this, we consider a deliberately simple setting in which the explanation provider strategically withholds explanations that expose the model’s reliance on sensitive attributes. We then evaluate, through systematic experiments, how difficult it is for an external auditor or observer to uncover this manipulation. In particular, we analyse detection under

various information settings, reflecting different levels of background knowledge available to the auditor.

Our findings show that even in this minimal setup, detecting cherry-picking is surprisingly difficult. These results have direct implications for regulatory frameworks and auditing practices. We conclude by discussing how auditors and policymakers can use these insights to design more robust mechanisms for transparency, ensuring that explainability serves its intended purpose rather than becoming a tool for obfuscation.

Our contributions are threefold:

- We formally define cherry-picking for counterfactual explanations.
- We empirically demonstrate that detecting cherry-picked counterfactuals is extremely difficult in practice.
- We provide practical guidance for explanation providers, auditors, and counterfactual explanation library developers on mitigating the risks associated with cherry-picking.

## References

- [1] Amina Adadi and Mohammed Berrada. “Peeking inside the black-box: a survey on explainable artificial intelligence (XAI)”. In: *IEEE access* 6 (2018), pp. 52138–52160.
- [2] Alejandro Barredo Arrieta et al. “Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI”. In: *Information fusion* 58 (2020), pp. 82–115.
- [3] Dieter Brughmans, Lissa Melis, and David Martens. “Disagreement amongst counterfactual explanations: how transparency can be misleading”. In: *Top* 32.3 (2024), pp. 429–462.
- [4] Bryce Goodman and Seth Flaxman. “European Union regulations on algorithmic decision-making and a “right to explanation””. In: *AI magazine* 38.3 (2017), pp. 50–57.
- [5] Satyapriya Krishna et al. “The Disagreement Problem in Explainable Machine Learning: A Practitioner’s Perspective”. In: *Transactions on machine learning research (TMLR)* (2024).
- [6] Wojciech Samek et al. *Explainable AI: interpreting, explaining and visualizing deep learning*. Vol. 11700. Springer Nature, 2019.
- [7] Feiyu Xu et al. “Explainable AI: A brief survey on history, research areas, approaches and challenges”. In: *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II* 8. Springer. 2019, pp. 563–574.

# Consistency in Persona-assigned Large Language Models

Manon Reusens

KU Leuven,

Universiteit Antwerpen

Bart Baesens

KU Leuven,

University of Southampton

David Jurgens

University of Michigan

**Keywords:** Persona-assigned Large Language Models, Consistency

Large Language Models (LLMs) are increasingly deployed across a wide range of applications, where assigning a persona allows tailoring model outputs to specific contexts. Although persona-based prompting has become a rapidly expanding area within Natural Language Processing, systematic investigations into the consistency of persona-driven responses remain limited. In particular, it is not yet clear to what extent LLMs maintain stable behavior within an assigned persona across diverse task types. Prior research has shown that LLM outputs are highly sensitive to prompt design; thus, variations in question format, e.g. multiple-choice versus open-ended prompts, may substantially influence the generated responses. Furthermore, little attention has been given to whether LLMs exhibit consistency not only within but also across personas that were not explicitly assigned.

To address these gaps, we introduce a novel framework for evaluating consistency in persona-assigned LLMs. Specifically, we assess both intra-persona consistency, the degree to which the model remains aligned with its assigned persona, and inter-persona consistency, the extent to which the model answers consistently across non-assigned personas. Grounded in social science literature, our framework evaluates consistency across five output dimensions: survey answering, social media post generation, essay writing, single-turn question answering, and multi-turn question answering. Our findings indicate that consistency is shaped by three key factors. First, models generally display stable behavior within their assigned persona. Second, inter-persona consistency appears strongly influenced by existing stereotypes associated with those personas. Finally, we observe evidence of model default values, where LLMs tend to produce socially desirable answers.

## Leveraging probabilistic imbalance price forecasts for risk-aware trading with battery systems

Margarida Mascarenhas

KU Leuven, ELECTA

e-mail: [margarida.mascarenhas@kuleuven.be](mailto:margarida.mascarenhas@kuleuven.be)

Arnaud Verstraeten

KU Leuven, ELECTA

e-mail: [arnaud.verstraeten@kuleuven.be](mailto:arnaud.verstraeten@kuleuven.be)

Hussain Kazmi

KU Leuven, ELECTA

e-mail: [hussain.kazmi@kuleuven.be](mailto:hussain.kazmi@kuleuven.be)

The energy transition has accelerated the rollout of variable renewable energy (VRE) generation technologies [1]. The electricity generation profiles of these technologies can be difficult to predict, which has complicated the delicate task of balancing electricity supply and demand. Individual balance responsible parties (BRPs) are incentivized to mitigate the system imbalance through the imbalance price, which is used to penalize or remunerate imbalanced positions. In the last five years, imbalance prices have increased and become more volatile in many European markets, indicating a higher required cost to balance the power system. Consequently, short-term imbalance price forecasting is gaining strategic importance for BRPs seeking to manage balancing costs [2], and for the transmission system operator (TSO) aiming to maintain system stability at minimal cost. To help mitigate this instability, our work develops a forecasting and control framework that combines data-driven imbalance price prediction with battery scheduling in the imbalance market.

We first developed a deterministic imbalance price forecast with a two hour horizon using an ensemble architecture, combining deep neural networks [3], regularized linear models [4], and gradient-boosted ensembles of decision trees [5]. After a sensitivity analysis, we observed that lagged imbalance prices, the day-ahead price, and the system imbalance are the most important input features. The ensemble is benchmarked against two naive methods and a trial forecast from Elia, the Belgian TSO. Our model achieves a 13% reduction in mean absolute error compared to Elia's forecast, and improves the naive baselines by 2-17%.

To quantify the operational value of this point forecast, we implement a deterministic receding horizon control strategy for a grid-connected battery participating in the imbalance market. The strategy uses the ensemble forecast to

select charge–discharge actions that maximise expected revenue while respecting operational constraints. To keep track of the state of charge, we split the trading period in blocks of two hours, and force the battery to have the same energy levels at the start and end of each block. This strategy leads to an annual operational profit of 5.7 EUR/MW/QH and chooses actions that mitigate the system imbalance 60% of the time, compared to the theoretical 75% upper bound achievable under perfect foresight.

However, imbalance prices are characterised by heavy tails, volatility, and spikes driven by real-time system events under which point forecasts may lead to suboptimal or overly aggressive control decisions. To model this uncertainty, we extend the ensemble framework using quantile regression to produce probabilistic imbalance price forecasts. These forecasts produce well-calibrated prediction intervals that represent the underlying uncertainty in the imbalance price distribution that point forecasts overlook.

Finally, we will integrate these probabilistic forecasts into a risk-aware battery optimisation scheme. Instead of maximising expected profit under a single trajectory, the controller incorporates risk measures to hedge against adverse price realisations. This enables more conservative charging during uncertain high-price intervals, or more opportunistic discharging when the upside potential is high but not guaranteed. Such a strategy is particularly valuable in the Belgian imbalance market, where penalties for wrong-way actions during extreme events can outweigh expected gains from average-case behaviour. By embedding quantile information directly into the optimisation problem, the battery can maximize profit in a more robust way, avoiding high-risk trades during periods with high forecast uncertainty.

The combination of probabilistic forecasting and risk-aware control provides a framework for operating battery energy storage systems in volatile balancing markets. By including additional constraints (e.g., thermal comfort) in the optimisation problem, the results can also be generalised to steer other flexible assets such as heat pumps or cooling installations. Given the inherent uncertainty of imbalance prices, quantile-based strategies allow BRPs to make more informed decisions and better manage exposure to adverse market outcomes.

## References

- [1] International Energy Agency, “World energy outlook 2025,” Paris, 2025, licence: CC BY 4.0 (report); CC BY NC SA 4.0 (Annex A). [Online]. Available: <https://www.iea.org/reports/world-energy-outlook-2025>
- [2] R. Smets, K. Bruninx, J. Bottieau, J.-F. Toubreau, and E. Delarue, “Strategic Implicit Balancing With Energy Storage Systems via Stochastic Model Predictive Control,” *Policy and Regulation IEEE Transactions on Energy Markets*, vol. 1, no. 4, pp. 373–385, Dec. 2023.
- [3] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, “Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark,” *Applied Energy*, vol. 293, p. 116983, Jul. 2021.
- [4] B. Uniejewski, J. Nowotarski, and R. Weron, “Automated Variable Selection and Shrinkage for Day-Ahead Electricity Price Forecasting,” *Energies*, vol. 9, no. 8, p. 621, Aug. 2016.
- [5] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” *CoRR*, vol. abs/1603.02754, 2016.

## Decision-focused learning for optimal PV-Battery scheduling

Joris Depoortere

KU Leuven, ESAT Electa

e-mail: [joris.depoortere@kuleuven.be](mailto:joris.depoortere@kuleuven.be)

Hussain Syed Kazmi

KU Leuven, ESAT Electa

e-mail: [hussain.kazmi@kuleuven.be](mailto:hussain.kazmi@kuleuven.be)

Johan Driesen

KU Leuven, ESAT Electa

e-mail: [johan.driesen@kuleuven.be](mailto:johan.driesen@kuleuven.be)

The use of residential photovoltaics has increased dramatically in recent years. With battery systems becoming more affordable, the optimal operation of a photovoltaic-battery system can bring significant savings to households. Optimal control of these systems requires correct forecasts of the underlying parameters, such as photovoltaic power generation, to correctly know when to charge and discharge the battery, given how much energy was generated. Forecasting models have become increasingly accurate due to algorithmic advances and data availability, but forecast accuracy is typically measured in generic metrics, which might not align with the goals of the downstream application.

Therefore, we propose to integrate optimization and forecasting in a combined approach, training a photovoltaic power forecaster on the downstream cost savings from optimal control of a battery system using these forecasts. We apply this framework to 20 different dutch buildings over a span of multiple years. This methodology, called decision-focused learning, integrates optimization into the learning process. Our methodology consists of training a long short-term memory model which has a convex optimization layer attached to it. This convex optimization layer allows us to differentiate through the optimization problem, which makes it possible to train the network on the results of the optimization problem through backpropagation. We evaluate the performance of our model on the regret: the difference in the solution of the downstream optimization given the parameter values given by the forecast model versus the use of the true parameter values.

Our results show that coupling the forecast with the optimization problem provides value for the user, decreasing the relative cost by approximately 8% over a two-phase approach using a generic statistical metric to train the forecast model. The decision-focused photovoltaic energy predictions are more conservative compared to the predictions made by a model trained on the mean squared error, which makes sense given that over-predicting is more expensive than under-predicting, as off-take from the grid comes at higher prices than injecting into the grid. Finally, We also note that warm starting a decision-focused model gives better performance over immediate decision-focused forecasting and makes it more robust to uncertainties in other parameters in the optimization problem.

# Mitigating the Bullwhip Effect when Forecasting for Inventory Control

Joost F. van der Haar

KU Leuven, Research Center for Operations Management

e-mail: [joost.vanderhaar@kuleuven.be](mailto:joost.vanderhaar@kuleuven.be)

Myra Dumon

KU Leuven, Research Center for Operations Management

e-mail: [myra.dumone@student.kuleuven.be](mailto:myra.dumone@student.kuleuven.be)

Jente Van Belle

KU Leuven, Research Center for Information Systems Engineering

e-mail: [jente.vanbelle@kuleuven.be](mailto:jente.vanbelle@kuleuven.be)

Robert N. Boute

KU Leuven, Research Center for Operations Management

e-mail: [robert.boute@kuleuven.be](mailto:robert.boute@kuleuven.be)

Mitigating the bullwhip effect is a key strategy to maintain and increase supply chain profitability. Relevant literature is hence ample, but this literature has thus far focused on mitigating the bullwhip effect when using statistical forecasting models. We propose a new method to mitigate the bullwhip effect when using machine learning forecasting models, which are increasingly replacing statistical forecasting models in practice. Machine learning models are often able to produce substantially more accurate forecasts, yet recent literature [1] suggests that these more accurate forecasts can increase order variability and thereby induce the bullwhip effect. We address this trade-off by introducing a custom loss function that combines forecast accuracy with a bullwhip component.

The bullwhip effect refers to the amplification of demand variability from down- to upstream in the supply chain. This distortion of demand information often leads to inefficiencies such as excessive inventory, stockouts, and overall increased operational costs. It is induced by demand forecasting errors, order batching, price fluctuations and lead time variability. Understanding and mitigating this effect remains a critical challenge for both practitioners and researchers.

Forecasting research seeks to minimize demand forecasting errors, and it was long thought that minimizing these errors would also minimize the bullwhip effect. However, a recent study by [1] showed that the increase in accuracy from using machine learning methods that leverage explanatory variables can lead to a stronger bullwhip effect rather than a weaker one. The intuition behind these results is that more accurate forecasting models are more sensitive to deviations in the demand distribution, leading to higher forecast variability and consequently also a stronger bullwhip effect.

Nonetheless, the field of forecasting is increasingly embracing machine learning

methods. By learning over multiple SKUs at the same time, a technique known as global learning, machine learning forecasting models are able to extract and predict more fine-grained patterns by greatly increasing the available sample size at a relatively small cost in additional data heterogeneity. Avoiding machine learning methods to mitigate the bullwhip effect is hence undesirable.

In this paper, we bridge the gap between machine learning forecasting methods and bullwhip effect reduction by proposing an approach that simultaneously mitigates the bullwhip effect and optimizes forecast accuracy. More specifically, we train machine learning models to minimize a custom loss function that contains both a Mean Squared Scaled Error (MSSE) and Bullwhip Ratio (BR) component, and weigh the two terms using a scaling factor  $\alpha \in [0, 1]$  that can be chosen based on business needs.

We evaluate our method on daily sales data from the M5-dataset using a standard prequential training/validation/test-split of 1913/28/28 days, the feature engineering of the M5-winning solution, LightGBM as choice for the supervised learning algorithm and Optuna for hyperparameter tuning. As daily forecasts only make sense for smooth and erratic time series from a business standpoint, we restrict ourselves to time series that can be categorized as either smooth or erratic according to the classification by [2]. A summary of key results can be found in Figure 1, which shows that our approach allows users to trade-off bullwhip effect mitigation and forecast accuracy according to their needs. Moreover, it shows that the bullwhip effect can be substantially mitigated at the cost of only a marginal loss in accuracy.

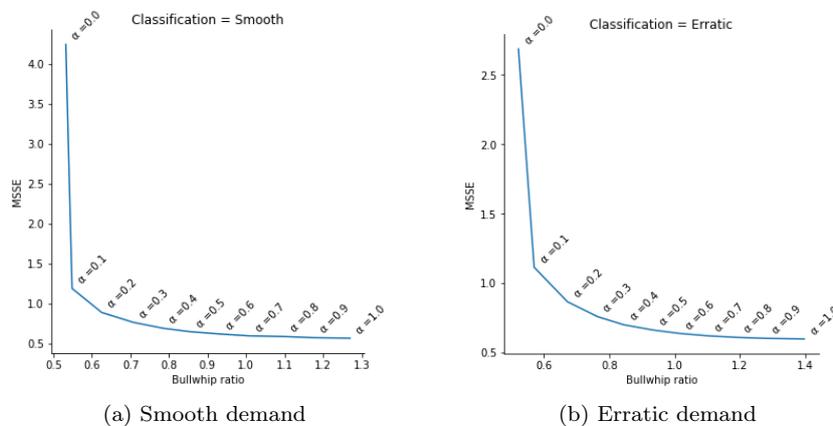


Figure 1: MSSE vs. BR for smooth and erratic demand under different  $\alpha$ -values

## References

- [1] Wellens, A.P., Boute, R.N., & Udenio Castro, M. (2023). Increased bullwhip in retail: A side effect of improving forecast accuracy with more data?. *Forecasting: The International Journal of Applied Forecasting*, 2023(68), 25-30.
- [2] Kostenko, A.V., & Hyndman, R.J. (2006). A note on the categorization of demand patterns. *Journal of the Operational Research Society*, 57(10), 1256-1257.

# Exploring Machine and Deep Learning Methods to Forecast Cash Flow in the Logistics Industry

Alice Wolfe

IESEG School of Management, Marketing and Sales

e-mail: [a.wolfe@ieseg.fr](mailto:a.wolfe@ieseg.fr)

Sarah Van der Auweraur

IESEG School of Management, Operations

e-mail: [s.vanderauweraer@ieseg.fr](mailto:s.vanderauweraer@ieseg.fr)

Kristof Coussement

IESEG School of Management, Marketing and Sales

e-mail: [k.coussement@ieseg.fr](mailto:k.coussement@ieseg.fr)

December 4, 2025

Company cash flow management is a critical operation that impacts daily operations, investments, revenue, and risk management. With enhanced cash flow forecasting, companies can improve their cash management, becoming more agile in their investment decisions and operations, which can bring significant increases in revenue [1].

Despite its high importance to any operating company, it remains a limited field of study within academic forecasting research and has focused largely on the utilization of traditional forecasting methods. Cash flow forecasting lags behind general financial forecasting in the exploration of Machine and Deep learning methods and their potential improvement to the forecast. The aim of this research is therefore to inspect the application of Machine Learning methods on company cash flow forecasting compared to traditional forecasting methods.

This research is carried out in close collaboration with an international logistics company in the fashion and lifestyle industry. The company provided a data set of approximately 8 million financial records consisting of client and supplier invoice and payment records from 2016 to 2025. This context-specific data is combined with factors related to contextual client and company data. Using this data, we evaluate traditional forecasting methods: Naive, ARIMA and ETS, with Machine Learning methods: Linear Regression, Support Vector Regression and Gradient Boosting, and Deep Learning methods: ANN and LSTM to provide a forecast of the company's cash flow across different time horizons [2] [3].

We implement a sliding window to evaluate the traditional methods and a blocked 5-fold cross validation method to evaluate the Machine and Deep Learning methods [5]. The models are evaluated using MSE to determine the best performing model and RMSE is used to bring the results back into the contextual units of the cash flow [4].

With our initial results we have found that when utilising only lagged values of the cash flow, the traditional methods have significantly exceeded the performance of the Naive benchmark. The machine learning methods do not improve on the traditional models, however with an LSTM model that utilizes temporal elements, we see a significant improvement over the traditional methods. When including additional features the results show that as the model increases in complexity, so does the forecast performance.

The Post-modeling analysis is completed using SHAP to explore the impact of additional features on the forecast and to assist in providing information to inform strategic decisions made by the company [3].

## References

- [1] M. Quails, “Why an Emphasis on Cash Flow Forecasting Remains Critical,” *The Wall Street Journal*, 2023.
- [2] H. Hassani, E. Silva, R. Gupta, and M. Segnon, “Forecasting the price of gold,” in *Applied Economics*, vol. 47, no. 39, pp. 4141–4152, 2015.
- [3] R. Tangsuecheeva and V. Prahbu, “Can Third-Party Brand Value Estimates Help Predict Cash Flows? A Machine-Learning Analysis,” in *European Accounting Review*, vol. 158, pp. 65–76, 2025.
- [4] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, OTexts, 2018.
- [5] V. Cerqueira, L. Torgo, and I. Mozetič, “Evaluating time series forecasting models: an empirical study on performance estimation methods,” *Machine Learning*, vol. 109, pp. 1997–2028, 2020. <https://doi.org/10.1007/s10994-020-05910-7>

## Enhancing population localization by geo-referencing buildings to census grids

Alexandre Kerff

University of Liège, HEC Liège Management School,  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [akerff@uliege.be](mailto:akerff@uliege.be)

Jan Weymeirsch

Trier University,  
Economic and Social Statistics Department  
e-mail: [weymeirsch@uni-trier.de](mailto:weymeirsch@uni-trier.de)

Morgane Dumont

University of Liège, HEC Liège Management School,  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [morgane.dumont@uliege.be](mailto:morgane.dumont@uliege.be)

Élise Vandomme

University of Liège, HEC Liège Management School,  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [elise.vandomme@uliege.be](mailto:elise.vandomme@uliege.be)

The geo-referencing of buildings and their inhabitants at detailed geographical levels can be required for multiple applications, for example, transportation planning, evaluation of housing policies, provision of health services, or flooding information. Indeed, the precise location of residence as well as place of employment of the individuals and households is crucial information in these fields. Spatial dynamic microsimulations especially project these individual units, over time and spatial dimensions, instead of aggregates to better understand behaviour at the individual level, when histories and heterogeneity are crucial [1]. The German model MikroSim [2], in particular, represents the over 80 million German inhabitants and their households as a fully synthetic statistical twin and directly benefits from a detailed representation of their location.

Individuals and household data often comes from census datasets, which depict aggregate population information in a grid cell representation of various resolutions. On the other hand, micro-level buildings datasets, such as ALKIS [3], represent individual buildings, their type, associated geospatial attributes and locations. The merging of these two sources of geographical information consists in a new task, the assignment of building objects to grid cells (Figure 1). The allocation of buildings to the census grid cells is neither fully transparent, nor

conducted purely on geographic localization alone, meaning that buildings can in principle be attributed to a neighbouring cell than that they are geographically located in. This problem can be represented as a bipartite graph, where each building must then be assigned to close grid cells, ensuring that the total estimated population of assigned buildings does not exceed the grid cell's capacity.

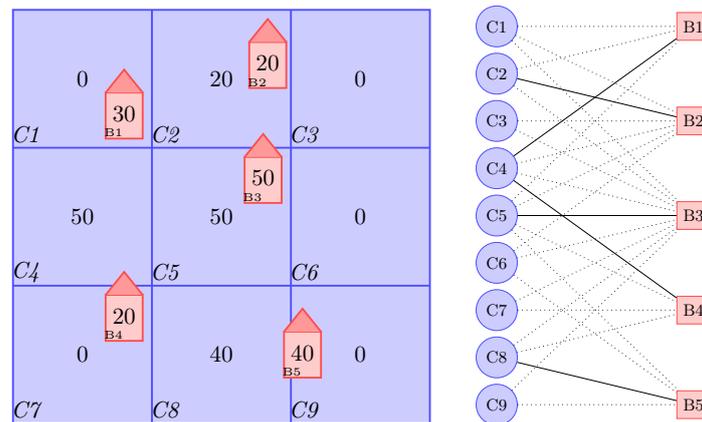


Figure 1: Grid cell assignment of buildings

This problem is similar to the class of generalized assignment problems (GAP) [4]. In this study, we present different mathematical formulations of this problem that aim at minimising differences in estimated building inhabitants and grid cells' population level; and minimising distances between buildings and assigned grid cells. This problem presents some challenges, such as allocation of building complexes, heterogeneous building types or uncertainty in buildings capacity estimates as well as potential register errors.

## References

- [1] Spielauer, M. (2011). *What is social science microsimulation?* Social Science Computer Review, 29(1), 9-20.
- [2] Münnich R., Schnell R., Brenzel H., Dieckmann H., Dräger S., Emmenegger J., Höcker P., Kopp J., Merkle H., Neufang K., Obersneider M., Reinhold J., Schaller J., Schmaus S., Stein P. (2021). *A population based regional dynamic microsimulation of Germany: the MikroSim model*. Methods, data, analyses: a journal for quantitative methods and survey methodology, 15(2), 241-264.
- [3] Working Committee of the Surveying Authorities of the Länder of the Federal Republic of Germany (AdV). *Authoritative Real Estate Cadastre Information System (ALKIS)*. Retrieved from <https://www.adv-online.de/Products/Real-Estate-Cadastre/ALKIS/>
- [4] Kundakcioglu, O.E. and Alizamir, S. (2008). *Generalized Assignment Problem*. In: Floudas, C. and Pardalos, P. (eds) *Encyclopedia of Optimization*. Springer, Boston, MA.

## Improving Pickup and Delivery Routing in a Laundry Service Company

Charlotte Moonen

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [charlotte.moonen@uliege.be](mailto:charlotte.moonen@uliege.be)

Cristian Aguayo

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [caguayo@uliege.be](mailto:caguayo@uliege.be)

Sabine Limbourg

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [sabine.limbourg@uliege.be](mailto:sabine.limbourg@uliege.be)

Household laundry represents a significant environmental burden, with residential washing machines accounting for substantial water and energy consumption across Europe. Professional laundry services offer a sustainable alternative by leveraging economies of scale and industrial efficiency, yet their environmental and operational benefits depend critically on optimized distribution logistics. Despite growing interest in sustainable service models, the literature remains sparse on transportation improvement strategies for laundry service operations, particularly for systems requiring simultaneous pickup and delivery with strict time windows.

This research examines the improvement of logistics operations for a Belgian laundry service company that manages both the execution of laundry services and the associated transportation flows. More specifically, it is responsible for the pickup of dirty laundry from customers and the delivery of clean laundry to individuals. The company faces two objectives: expanding its customer base while guaranteeing a maximum 72-hour turnaround time between collection and delivery. These objectives necessitate a fundamental redesign of the distribution network to maintain service quality while controlling costs and minimizing environmental impact through reduced vehicle kilometers traveled.

The problem is formulated as a Vehicle Routing Problem with Simultaneous Pickup and Delivery (VRPSPD), specifically a multiple-vehicle Hamiltonian one-to-many-to-one Pickup-and-Delivery Problem with combined demands [1]. This formulation addresses a critical gap in the routing literature, as VRPSPD applications remain underexplored in service industries despite their practical relevance. The company must design daily routes starting and ending at a central depot, serving customers distributed throughout Wallonia who may require pickup, de-

livery, or both. The objective is to determine vehicle routes that minimize total travel distance under the operational constraints imposed by the service process.

To address this routing problem, several solution approaches have been applied. First, given the spatial nature of VRP-type problems, Geographic Information Systems (GIS) were integrated with OR techniques to visualize routing scenarios [2]. Based on the company's initial solution, manual improvements were performed using cross-route customer exchanges inspired by 2-opt moves extended to multiple routes [3]. Second, a mixed-integer programming model was developed to represent the VRPSPD more rigorously and solved using the Gurobi MIP solver [4].

Beyond solving the current VRPSPD, the research addresses a broader operational challenge since the company's logistics model is likely to evolve. Alternative distribution structures such as the integration of drop-off points or parcel lockers, which are increasingly widespread, are currently under study. Adopting such models would require the development of new modeling formulations and solution methods tailored to the organization's future logistics environment.

### Acknowledgment

The authors gratefully acknowledge the support provided by the Walloon Region for funding this research through the LaaS project (grant number C-9095), as well as the Walloon clusters Mecatech and Logistics in Wallonia.

### References

- [1] Çağrı Koç, Gilbert Laporte, İlknur Tükenmez, *A review of vehicle routing with simultaneous pickup and delivery*, *Computers & Operations Research*, 2020, vol. 122, p. 104987. <https://www.sciencedirect.com/science/article/pii/S0305054820301040>
- [2] Saoussen Krichen, Sami Faiz, Takwa Tlili, Khaoula Tej, *Tabu-based GIS for solving the vehicle routing problem*, *Expert Systems with Applications*, 2014, vol. 41, no. 14, pp. 6483-6493. <https://www.sciencedirect.com/science/article/pii/S0957417414001651>
- [3] Christian Prins, *A simple and effective evolutionary algorithm for the vehicle routing problem*, *Computers & Operations Research*, 2004, vol. 31, no. 12, pp. 1985-2002. <https://www.sciencedirect.com/science/article/pii/S0305054803001588>
- [4] Alan Osorio-Mora, Mauricio Soto-Bustos, Gustavo Gatica, Pedro Palominos, Rodrigo Linfati, *The Multi-Depot Cumulative Vehicle Routing Problem With Mandatory Visit Times and Minimum Delayed Latency*, *IEEE Access*, 2021, vol. 9, pp. 27210-27225. <https://doi.org/10.1109/ACCESS.2021.3058242>

## Evolution of social acceptability of public lighting reductions throughout the recent energy crisis

Elodie Bebronne

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [elodie.bebonne@uliege.be](mailto:elodie.bebonne@uliege.be)

Sabine Limbourg

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative Methods and Operations Management  
e-mail: [sabine.limbouurg@uliege.be](mailto:sabine.limbouurg@uliege.be)

Mario Cools

University of Liège, Faculty of Applied Sciences  
Urban and Environmental Engineering - LEMA  
e-mail: [mario.cools@uliege.be](mailto:mario.cools@uliege.be)

Artificial light at night (ALAN) is among the most important technological developments. It has facilitated nocturnal human activity, extending daytime activities into the night. However, ALAN is increasingly recognised as an environmental concern due to its contribution to light pollution, which adversely impacts both human health and wildlife [1, 2]. This duality presents a complex challenge: how can public spaces be illuminated in an economically, socially and environmentally sustainable manner?

Despite increasing interest in reducing public lighting to mitigate environmental impacts and energy consumption, the success of such strategies depends on political commitment, which may be limited by expected public opposition [3]. There is a lack of repeated cross-sectional evidence capturing changes in public acceptability towards lighting reduction over time.

The present research is original in its capture of the evolution of social acceptability over time, throughout the recent energy crisis. We conducted a regional cross-sectional study of the Walloon population. Overall, we collected 3 114 responses across three surveys distributed before the energy crisis escalation (February-March 2022), at its peak (October-November 2022), and two years after (October-November 2024). Data analysis employs binary logistic regression and inferential statistics. The study examines the association between acceptability and socio-demographic, contextual, and perceptual factors, including safety concerns and mobility in public spaces. Public opinion is more positive than local political authorities may have expected, with support increasing during the crisis and remaining stable afterwards.

**Keywords:** Social acceptability, surveys, binary logistic regression, public lighting, sustainability, light pollution, nocturnal mobility.

## References

- [1] Ron Chepesiuk. Missing the dark: health effects of light pollution. *Environmental Health Perspectives*, 2009, 117:A20–A27. doi.org/10.1289/ehp.117-a20.
- [2] Kevin Gaston, and Alejandro Sánchez de Miguel. Environmental impacts of artificial light at night. *Annual Review of Environment and Resources*, 2022, 47:373–398. doi.org/10.1146/annurev-environ-112420-014438.
- [3] Beaudet, Chloé, Léa Tardieu, and Maia David. Are citizens willing to accept changes in public lighting for biodiversity conservation? *Ecological Economics*, 2022, 200:107527.

# Stochastic Revelation in Sequential Decision-Making with a Pharmaceutical Application

Morteza Davari

SKEMA Business School,  
Center of Analytics and Management Science  
e-mail: [morteza.davari@skema.edu](mailto:morteza.davari@skema.edu)

Merve Bodur

The University of Edinburgh,  
School of Mathematics  
e-mail: [merve.bodur@ed.ac.uk](mailto:merve.bodur@ed.ac.uk)

Bernardo Pagnoncelli

SKEMA Business School  
Center of Analytics and Management Science  
e-mail: [bernardo.pagnoncelli@skema.edu](mailto:bernardo.pagnoncelli@skema.edu)

**Keywords:** Stochastic programming, sequential decision-making, uncertainty revelation

In many real-life situations, such as medical product launches, energy investments, or the rollout of new policies, decision-makers must act before knowing exactly when critical information will become available. We develop new mathematical models that incorporate uncertainty about what will happen and when that uncertainty will resolve. Traditional decision-making tools assume fixed timelines for information revelation; instead, we address more realistic scenarios where information arrives at unpredictable moments, making planning more complex and costly if mishandled.

In this work, we focus on a pharmaceutical application, in particular, the planning and production of medical devices, where firms must be mindful of regulatory approvals. This context has long been recognized as a challenge due to uncertain approval dates, which complicate inventory and production planning decisions [1]. The change in demand resulting from regulatory uncertainty is closely connected to that of optimal product rollover, where the timing of phasing out old products and introducing new ones significantly affects demand and, as such, the firm's performance [2].

Our first contribution is to formulate the problem as a multistage stochastic programming model [3, 4]. Moderate-sized instances can be solved using the extensive form, which provides insights into both the solution structure and the design of efficient algorithms. We then propose decomposition methods that enable the solution of larger instances, such as those involving more countries and longer planning horizons, and explore the managerial insights derived from these solutions. Finally, we develop a branch-and-bound approach based on dominance rules for production-switch decisions, together with a lower bound obtained from the linear relaxation of the full formulation, which can be formulated as a Marko-

vian SDDP, to further improve the computational tractability of the problem.

Our preliminary results show that the branch-and-bound approach greatly outperforms the decomposition methods and the full extensive-form model. However, the non-scalability of the scenario tree leads to many instances that remain unsolved, particularly when the number of countries exceeds ten.

## References

- [1] A. V. Hill and W. J. Sawaya, Production planning for medical devices with an uncertain regulatory approval date, *IIE Transactions*, 36(4): 307–317, 2004.
- [2] W. S. Lim and C. S. Tang, Optimal product rollover strategies, *European Journal of Operational Research*, 174(2): 905–922, 2006.
- [3] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, Springer, 2nd edition, 2011.
- [4] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, SIAM, 2nd edition, 2014.

## Optimizing Urban Logistics by Integrating a Construction Consolidation Center

Yunyun Wang

Hasselt University, Research Group Logistics

e-mail: [yunyun.wang@uhasselt.be](mailto:yunyun.wang@uhasselt.be)

An Caris

Hasselt University, Research Group Logistics

[an.caris@uhasselt.be](mailto:an.caris@uhasselt.be)

Kris Braekers

Hasselt University, Research Group Logistics

[kris.braekers@uhasselt.be](mailto:kris.braekers@uhasselt.be)

Ruben D'Haen

Hasselt University, Research Group Logistics

[ruben.dhaen@uhasselt.be](mailto:ruben.dhaen@uhasselt.be)

This study investigates the integration of a Construction Consolidation Center (CCC) into urban construction logistics. A CCC functions as a consolidation and distribution hub, enabling the bundling of materials and facilitating Just-In-Time (JIT) deliveries [Guerlain et al., 2019]. The overall objective of our research is to enhance the efficiency of urban construction logistics while minimizing costs and social burdens. Various scenarios and operational parameters are analyzed to quantify the benefits of a CCC for both public and private stakeholders, including the potential for reduced urban traffic and cost savings for construction companies and material suppliers.

Our problem setting involves multiple construction sites located within an urban area, which requires materials to be supplied from external locations. The CCC, typically located at the outskirts of the city, offers storage facilities and a dedicated vehicle fleet to ensure timely delivery to construction sites. Construction sites, constrained by limited on-site storage, request materials according to their project schedules, creating a deterministic demand. Suppliers determine the amount of materials to be delivered directly to construction sites or to the CCC, where the materials are stored before being dispatched to their final destinations.

We propose a mathematical model to optimize the daily delivery plans for both suppliers and the CCC. Decisions include vehicle assignment, scheduling, and route planning for two types of deliveries: (1) deliveries from suppliers directly to construction sites or to the CCC and (2) deliveries from the CCC to the construction sites. The objective of the model is to minimize total costs, which comprise

transportation costs for suppliers and the CCC, handling costs at the CCC, as well as inventory costs incurred at the CCC and construction sites.

A rolling horizon approach is adopted, inspired by [Nolz, 2021], to accommodate disturbances and unforeseen changes. Daily plans are created for an extended horizon, with the first week's plan being fixed, while subsequent plans are updated on a weekly basis.

To solve this optimization model, a two-phase solution algorithm is developed. In the first phase, we determine the mode (direct delivery or via CCC) and the delivery timing for each request. At this stage, vehicle assignment and routing constraints are relaxed to approximate transportation costs, allowing for an optimized total cost. In the second phase, further vehicle assignment and route planning are established using a heuristic algorithm, with optimization at both the daily level and across multiple days.

Experimental results are presented, comparing a baseline scenario against various operational configurations. We analyze the specific impact of CCC integration and on-site time windows, supported by a sensitivity analysis of key parameters. Ultimately, this study provides critical insights into the operational and policy implications of CCC adoption, offering a pathway for sustainable and efficient urban construction logistics.

## References

- C. Guerlain, S. Renault, and F. Ferrero, "Understanding Construction Logistics in Urban Areas and Lowering Its Environmental Impact: A Focus on Construction Consolidation Centres," *Sustainability*, vol. 11, no. 21, p. 6118, Nov. 2019. Available: <https://www.mdpi.com/2071-1050/11/21/6118>. doi: 10.3390/su11216118.
- P. C. Nolz, "Optimizing Construction Schedules and Material Deliveries in City Logistics: A Case Study from the Building Industry," *Flexible Services and Manufacturing Journal*, vol. 33, no. 3, pp. 846–878, Sep. 2021. Available: <https://link.springer.com/10.1007/s10696-020-09391-7>. doi: 10.1007/s10696-020-09391-7.

# An application of advanced optimization methods in Non-Emergency Medical Transportation

Bryan Galarza Montenegro

University of Antwerp, Department of Engineering Management (ENM)

ANT/OR - Antwerp Operations Research Group

e-mail: [bryan.galarzamontenegro@uantwerpen.be](mailto:bryan.galarzamontenegro@uantwerpen.be)

Healthcare systems are increasingly shifting toward shorter hospital stays and outpatient care, with discharge rates rising across much of the EU [1]. As more treatments occur outside hospitals, demand for ambulatory and home-based services continues to grow. Combined with demographic aging, this trend is driving a sharp increase in medical transportation needs, not only for patients, but also for staff, equipment, pharmaceuticals, and diagnostic samples [2].

Non-Emergency Medical Transportation (NEMT) represents a key solution to these accessibility challenges, particularly for people who require regular access to medical care but do not need emergency services [3]. NEMT provides transportation for patients and/or medical equipment to routine medical appointments, such as rehabilitation visits, dialysis, sample delivery and check-ups. It plays a vital role in reducing missed appointments and ensuring timely healthcare access. Despite their importance, current NEMT systems often face challenges related to inefficiency, long waiting times, and poor coordination with healthcare providers.

NEMT planning is operationally challenging. In practice, dispatchers manually assign trips to a heterogeneous fleet of vehicles and staff shifts. Furthermore, they handle complex medical and operational constraints, and adapt to dynamic conditions. This planning often requires several hours and depends on expert knowledge. As part of a VLAIO development project, I designed and implemented an optimization algorithm to automate this process. The algorithm is embedded in the dispatching platform of Triptomatic [4], a prominent software company in the sector.

The optimization algorithm receives patient trips (with tasks, time windows, mobility requirements), staff shifts, heterogeneous vehicles, and a road network. Trips may include multiple pick-up and drop-off tasks. Vehicles vary in capacity and capabilities (e.g. normal seats, wheelchair seats, stretcher spaces), medical equipment/staff, and start/end locations. Key constraints include a.o. the assignment of trips to valid vehicle(shifts), routing and scheduling under soft and hard time windows, and combination of compatible trips into shared rides. The optimization algorithm also deals with case-specific constraints, such as the black-listing of trips to certain shifts. The objective is to minimize total operational

cost while maintaining service quality. A weighted, normalized cost function aggregates components such as lateness, overtime, empty driving, unassigned trips, and constraint-specific penalties.

The optimization algorithm is implemented in C++ and hosted on a remote server. It is embedded in the webapp of Triptomatic and directly interacts with the Triptomatic database. The core algorithm is a parallelized Large Neighborhood Search (LNS) metaheuristic. Each iteration destroys parts of the solution (removing assigned trips or shifts) and repairs them using randomized greedy heuristics. Simulated Annealing with adaptive cooling prevents premature convergence, while local search algorithms (including 2-opt) refine routes. Separate algorithmic modules handle case-specific restrictions, such as prescheduled trips, sample delivery routes and time-on-location calculations. Parallel execution across multiple threads accelerates convergence and improves solution robustness.

The engine supports both static (full-day planning) and real-time scheduling. Real-time optimization incorporates newly added or cancelled trips, dynamic vehicle positions, delay propagation, and dispatcher overrides. “Context trips” preserve continuity of existing schedules, allowing the system to blend automation with human decision-making.

The optimization engine was validated with real operational data from transport providers in Belgium and Germany. Dispatchers reported reduced workload and improved visibility into operations. Extensive unit testing ensured feasibility preservation across updates, and iterative feedback from practitioners guided refinements. The engine is now fully deployed in production.

This work demonstrates the successful application of advanced optimization methods to real-world NEMT operations. The integrated engine significantly improves planning efficiency and service quality, and its deployment in production environments highlights both its practical value and its potential as a basis for future fundamental research in transportation and healthcare logistics.

## References

- [1] Eurostat. (2024) Hospital discharges and length of stay statistics. Accessed: 2025-12-05. [Online]. Available: [https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Hospital\\_discharges\\_and\\_length\\_of\\_stay\\_statistics](https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Hospital_discharges_and_length_of_stay_statistics)
- [2] J. Morrissey, “Companies respond to an urgent health care need: Transportation,” *The New York Times*, 2018.
- [3] J. Sedlack and A. Aamer, “Examining non-emergency medical transportation in Maine: Challenges and opportunities in healthcare logistics”, *Global Journal of Flexible Systems Management*, vol. 26, pp. 529–542, 2025.
- [4] B. D. Galarza Montenegro. (2025) Triptomatic: AI-powered dispatching. Accessed: 2025-12-05. [Online]. Available: <https://triptomatic.com/en-be/nemt#dispatching>.

# A Large Neighborhood Search Approach for Large-Scale Fleet Assignment Problems

Reinout Annaert, Pieter Smet, Greet Vanden Berghe

KU Leuven, Numerical Analysis and Applied Mathematics (NUMA)

e-mail: {reinout.annaert, pieter.smet, greet.vanden.berghe}@kuleuven.be

The global market for vehicle and equipment rental, hereafter referred to as resource rental, has grown considerably in recent years (Market Reports World, 2024). This growth is driven by rising demand for flexible access to resources and a shift away from the ownership of expensive assets, with companies and individuals instead renting them as needed. The resource-rental providers themselves also benefit, as they can rely on recurring revenue streams rather than more volatile one-time sales. However, renting rather than selling resources also introduces unique operational challenges. Fleet assignment is a practical problem that every rental company encounters on a daily basis: providers must determine which rental requests to accept and which resources to assign to them. They must also decide when it is worthwhile to offer a customer a substitution if the requested category is unavailable, and whether transfers between depots are justified given their cost and the resulting temporary unavailability of the resource. These questions form the core of the fleet assignment problem considered in this study.

The efficiency of fleet assignment algorithms is not only important for long overnight planning runs, but also for the inherently dynamic nature of day-to-day operations. When a customer calls, there is no time to run an algorithm for minutes or hours before deciding whether the rental request can be accepted. Likewise, if a resource breaks down or a rental request unexpectedly overruns, operators must be able to generate an updated schedule very quickly.

Fleet assignment problems have been addressed using both synthetic and real-world data, under various assumptions and at different scales. The tractability of exact approaches (Faneyte et al., 2002; Ernst et al., 2010) depends not only on problem scale but also on resource-specific constraints, the number of unique locations, and the time granularity. Heuristic approaches (Ernst et al., 2011; Oliveira et al., 2014; Kulkarni et al., 2018) are often based on these exact approaches or incorporate them as components, and are therefore similarly affected by these factors. Both exact and heuristic approaches therefore struggle to achieve the scale and complexity demanded by many real-world problems.

While existing approaches can provide good schedules for small and medium-sized instances, they often require significant computation time before an initial feasible schedule is available. To address this challenge, this study proposes a guided ruin and recreate (GRR) heuristic. By integrating the GRR heuristic with metaheuristics such as Late-Acceptance Hill Climbing (LAHC) (Burke and Bykov, 2017) and Iterated Local Search (ILS) (Lourenço et al., 2003), the GRR

heuristic benefits from both diversification strategies. To improve scalability, we employ exact decomposition techniques and introduce a problem-agnostic workload distribution scheme that enables efficient parallel computing.

The broad applicability of the approach is demonstrated through three case studies. The cases differ in terms of their problem scale, fleet composition, transfer operations and decision horizons. To support future research, datasets derived from these industrial cases are made publicly available.

## References

- Burke, E.K., Bykov, Y., 2017. The late acceptance hill-climbing heuristic. *European Journal of Operational Research* 258, 70–78. doi:<https://doi.org/10.1016/j.ejor.2016.07.012>.
- Ernst, A.T., Gavrilouk, E.O., Marquez, L.O., 2011. An efficient lagrangean heuristic for rental vehicle scheduling. *Computers & Operations Research* 38, 216–226. doi:<https://doi.org/10.1016/j.cor.2010.04.014>.
- Ernst, A.T., Horn, M., Kilby, P., Krishnamoorthy, M., 2010. Dynamic scheduling of recreational rental vehicles with revenue management extensions. *Journal of the Operational Research Society* 61, 1133–1143. doi:<https://doi.org/10.1057/jors.2009.78>.
- Faneyte, D., Spieksma, F., Woeginger, G., 2002. A branch-and-price algorithm for a hierarchical crew scheduling problem. *Naval Research Logistics* 49, 743–759. doi:<https://doi.org/10.1002/nav.10044>.
- Kulkarni, S., Patil, R., Krishnamoorthy, M., Ernst, A., Ranade, A., 2018. A new two-stage heuristic for the recreational vehicle scheduling problem. *Computers & Operations Research* 91, 59–78. doi:<https://doi.org/10.1016/j.cor.2017.11.004>.
- Lourenço, H.R., Martin, O.C., Stützle, T., 2003. Iterated local search, in: *Handbook of meta-heuristics*. Springer, pp. 320–353. doi:[https://doi.org/10.1007/0-306-48056-5\\_11](https://doi.org/10.1007/0-306-48056-5_11).
- Market Reports World, 2024. Vehicle rental market size, share, growth, and industry analysis, by type (economy cars, luxury cars, suvs, vans, electric vehicles), by application (tourism, business, local commute, event transportation), regional insights and forecast to 2033. URL: <https://www.marketreportsworld.com/market-reports/vehicle-rental-market-14719818>.
- Oliveira, B., Carravilla, M., Oliveira, J., Toledo, F., 2014. A relax-and-fix-based algorithm for the vehicle-reservation assignment problem in a car rental company. *European Journal of Operational Research* 237, 729–737. doi:<https://doi.org/10.1016/j.ejor.2014.02.018>.

## A Mixed Integer Programming model for the Strategic Assembly Line Feeding Problem

Hannah Verplancke

Ghent University, Department of Business Informatics and Operations Management

FlandersMake@UGent — core lab CVAMO, Ghent, Belgium

e-mail: [hannah.verplancke@ugent.be](mailto:hannah.verplancke@ugent.be)

Emmanouil Thanos

Ghent University, Department of Business Informatics and Operations Management

FlandersMake@UGent — core lab CVAMO, Ghent, Belgium

e-mail: [emmanouil.thanos@ugent.be](mailto:emmanouil.thanos@ugent.be)

Gohar Porbar

Ghent University, Department of Industrial Systems Engineering and Product Design

FlandersMake@UGent — core lab ISyE, Ghent, Belgium

e-mail: [gohar.porbar@ugent.be](mailto:gohar.porbar@ugent.be)

El-Houssaine Aghezzaf

Ghent University, Department of Industrial Systems Engineering and Product Design

FlandersMake@UGent — core lab ISyE, Ghent, Belgium

e-mail: [elhoussaine.aghezzaf@ugent.be](mailto:elhoussaine.aghezzaf@ugent.be)

Veronique Limère

Ghent University, Department of Business Informatics and Operations Management

FlandersMake@UGent — core lab CVAMO, Ghent, Belgium

e-mail: [veronique.limere@ugent.be](mailto:veronique.limere@ugent.be)

The industrial shift from mass production to mass customization has driven manufacturing companies to evolve towards high variety low volume mixed model assembly lines [2]. Consequently, their internal logistics systems must be designed to handle larger part variety and increased material flow complexity [1].

The Strategic Assembly Line Feeding Problem (SALFP) integrates long-term decisions on facility layout and vehicle fleet sizing with tactical decisions on part feeding policy assignment and route selection. Addressing those decision levels in an integrated manner allows for cost efficient internal logistics system design with respect to long-term operations.

In more detail, the following interrelated decisions are optimized within the

SALFP. With respect to facility layout, the number, positions and sizing of preparation cells are determined. Regarding vehicle fleet investments, the numbers of vehicles acquired from different vehicle types are determined. Furthermore, the selection of part feeding policies is included. Finally, the routes which will be employed to deliver parts to preparation areas and the border of line must be determined.

We propose a Mixed Integer Programming (MIP) model to solve the SALFP. Preparation cells can be positioned in dedicated areas at any available location of the initial facility layout. The model considers three different vehicle types, i.e., tow trains, forklifts and automated guided vehicles. The five main part feeding policies, line stocking, boxed supply, sequencing, stationary and traveling kitting, as identified by [3] are considered and routes can be selected out of a predetermined set. To provide good-quality, feasible solutions it is essential to ensure continuous stock availability. To this end, the model aligns consumption rates, preparation cell locations, available vehicle fleet, and reserved material space at the border of line in order to prevent line stoppages.

The proposed model is evaluated using a wide range of instances generated from real-world data. We analyze the impact of vehicle fleet, route selection and feeding policy decisions on total costs, and provide managerial insights with respect to designing efficient line-feeding systems.

## References

- [1] Battini, D., Boysen, N., and Emde, S. (2013). Just-in-time supermarkets for part supply in the automobile industry. *Journal of Management Control*, 24(2):209–217.
- [2] Emde, S. and Boysen, N. (2012). Optimally locating in-house logistics areas to facilitate jit-supply of mixed- model assembly lines. *International Journal of Production Economics*, 135(1):393–402.
- [3] Schmid, N. A. and Limère, V. (2019). A classification of tactical assembly line feeding problems. *International Journal of Production Research*, 57:7586–7609.

# Metaheuristics: beyond the horse race

Kenneth Sörensen

University of Antwerp, Department of Engineering Management

ANT/OR — University of Antwerp Operations Research Group

e-mail: `kenneth.sorensen@uantwerpen.be`

*Metaheuristics* research arose around the same time as research on exact methods, in the period following World War II. However, unlike exact algorithms, grounding metaheuristics in mathematical theory has proven largely fruitless. An alternative scientific paradigm for metaheuristics has been elusive, and metaheuristics researchers have resorted to turning the field into a “horse race”, where good algorithmic performance equates with good science. This has given us fantastic algorithms that power real large-scale optimization systems. However, there is still very little in the way of science in the field of metaheuristics. We argue that there *does* exist a better way to do science in the field of metaheuristics. In this talk, we argue that the way forward is not mathematics, but empirics.

In our recent meta-analysis [1], we applied rigorous statistical methods from medical research to evaluate the true impact of adaptiveness in Adaptive Large Neighborhood Search (ALNS). The findings revealed that the benefits of adaptiveness are considerably less than commonly assumed, highlighting how accepted “truths” in the field often rely on anecdotal evidence rather than solid data.

To rectify this, the field requires a robust framework for empirical research. As detailed in [2], this requires the adoption of standardized experimental protocols, open data repositories, and modular algorithmic components. However, the most critical change is cultural: as a research community, we must make the mental shift from merely seeking faster algorithms to actively pursuing verifiable scientific knowledge.

## References

- [1] Turkeš, R., Sörensen, K., & Hvattum, L. M., 2021, Meta-analysis of metaheuristics: Quantifying the effect of adaptiveness in adaptive large neighborhood search. *European Journal of Operational Research*, 292(2), 423-442.
- [2] Swan, J., et al., 2022, Metaheuristics “in the large”. *European Journal of Operational Research*, 297(2), 393-406.

# Locating a facility in space: A review of cover and facility location problems in space-based applications

Pieter Deleye

University of Antwerp, Department of Engineering Management

e-mail: [pieter.deleye@uantwerpen.be](mailto:pieter.deleye@uantwerpen.be)

Christof Defryn

University of Antwerp, Department of Engineering Management

e-mail: [christof.defryn@uantwerpen.be](mailto:christof.defryn@uantwerpen.be)

Over the past two decades, space-based applications have become deeply integrated into our lives by enabling, amongst others, global navigation, communication, and Earth observation. In support of these services, the number of satellites in orbit is growing fast. As a result, the risk of collisions and failures has gone up. In response, commercial ventures have recently started the offering of space-based services, including satellite refuelling and repositioning [1,2]. With the emergence of such on-orbit satellite servicing (OOS), fundamental operational questions arise, foremost among them the question of where to locate assets. These assets range from single servicing spacecraft to in-orbit depots equipped with shuttles to execute refuelling or repair operations. While location theory has long been studied in terrestrial contexts, the space environment poses its own, unique challenges: determining a position in space involves six degrees of freedom and all assets are constantly in motion.

This study presents the results of a systematic literature review on locating decisions in space-based applications based on 183 scientific publications and covering three large subfields: orbit selection, constellation design, and on-orbit servicing. We assess how classical Operations Research problems — such as coverage,  $n$ -cover,  $p$ -median, facility location, and location-routing — have been adapted to comply with the space context. Moreover, we discuss modelling assumptions, constraints, and solution techniques used to address the high-dimensional nature of space-based location problems.

The findings of this work are multifold. First, we find that applications drive modelling decisions and the employed solution method. Positioning applications mainly require time-continuous  $n$ -cover problem solving and are unbothered by capacity [4]. This is distinctly different from continuous cover problems with capacity constraints for communication constellations [3]. Second, the main focus in the literature is on cover problems rather than facility location problems. This can primarily be explained by the orbital mechanics governing satellites' trajectories and the use of radio signals, resulting in traditional (Euclidean) distance measures — typical for evaluating facility location decisions — holding less

information. Third, from a methodological perspective the main body of literature relies on parametric design rather than combinatorial optimisation. These methods aim for an optimal value of the six parameters that define an orbit in space without a predefined set of candidate locations from which a subset is to be opened. This observation directly leads to our final observation, namely, that solution space reduction and enumeration is the dominant strategy. This is done either by setting stringent solution feasibility requirements, or by substantial discretisation of the considered parametric values.

Furthermore, we conclude that the field on location analysis in space holds large potential with regards to robust optimisation. The current literature does not rely on a systematic use of well-founded Operations Research practices and, therefore, lacks transparency regarding modelling approaches and variable considerations, hindering reproducibility. Robustness and sensitivity analyses are only rarely included.

Our study serves as good benchmark for future research on the optimisation of locating decisions in space-based applications.

## References

- (1) Abhijeet Kumar. "Emerging Commercial Services and Capabilities Advancements in On-Orbit Satellite Servicing". In: *AIAA AVIATION FORUM AND ASCEND 2025*. 2025. DOI: 10.2514/6.2025-4107.
- (2) Roberto Opromolla et al. "Future in-orbit servicing operations in space traffic management context". In: *Acta Astronautica 220* (2024). pp. 469-477. ISSN: 0094-5765. DOI: <https://doi.org/10.1016/j.actaastro.2024.05.007>.
- (3) R. Wang, X. Ye, P. He, Y. Cui and D. Wu, "LEO Satellite Constellation Design for Seamless Global Coverage with QoS Guarantee," *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*, Kuala Lumpur, Malaysia, 2023, pp. 2329-2334, doi: 10.1109/GLOBECOM54140.2023.10437778.
- (4) Z.L. Deng et al. "Optimization design of two-layer Walker for LEO navigation augmentation using a dynamic multi-objective differential evolutionary algorithm based on elite guidance". In: *GPS SOLUTIONS 27* (2023). ISSN: 1080-5370, 1521-1886. DOI 10.1007/s10291-022-01366-5.

## What many operations researchers have done wrong and what are the remedies

Thomas Stützle

Université libre de Bruxelles and F.R.S.-FNRS

IRIDIA

e-mail: [stuetzle@ulb.ac.be](mailto:stuetzle@ulb.ac.be)

What many operations researchers have done wrong in their research is that they do not differentiate between training instances and an independent test set when testing their newly developed algorithms. Instead, they often train the algorithm on the same (or a subset of) instances and then report the good results of their algorithm in papers [1, 2, 5, 6].

This is a major problem in the operations research community.

A question arises: how strong is the error's role in optimization? The answer depends on how many instances are considered in the training process. To illustrate this, we present simple results for a WalkSAT algorithm with one parameter—the walk probability—on random 3-SAT problems with varying numbers of clauses. These results show that as we progress from 1 training instance to 10 training instances and then to 100 training instances, the performance generally becomes better and more stable.

Let us focus on the issue involving a finite or potentially infinite set of future instances. Examples include recurring problems like vehicle routing or scheduling for each working day. This is the most common scenario in operations research. We argue that there is little value in an algorithm which is trained and tested on the same instances you have provided.

A completely different situation arises when there is only one instance of interest, such as the wiring of telephone lines in a country. In such cases, there is no training set; instead, the best result for the instance is the optimal result achieved so far. Here, the goal is not to develop an algorithm that performs well across various instances but rather to achieve optimal performance on the specific instance.

Anything between these two scenarios can be discussed in terms of known and unknown instances.

How can this problem be addressed? In the first scenario, the most common one, there are several alternatives. When dealing with hand-crafted instances, random instances, or encodings of other instances, one can separate a training set from a test set. For benchmark sets, the benchmark instances can serve as the test set, while an independent training set is created. How can the independent test set be generated? One has various options regarding instance sizes, distributions of values, etc., and must choose these appropriately. This choice also depends

on whether the goal is to develop an algorithm that challenges existing ones. Alternatively, one might consider different distributions that the algorithm then can utilize.

For individual application instances, such as those from industrial companies, known techniques can be used to generate additional instances. One approach is to take a known instance and produce similar values. For example, if you have a processing time  $p$ , you can use a random distribution with processing times in the range  $[0, (1 - x)p), (1 + x)p]$ , where  $x \in [0, 1]$ , and apply this method to all processing times. Similar methods exist for generating both smaller and larger instances.

One question remains: how can we ensure that training instances differ from test instances? First, everyone should be aware of the error and adopt proper practices, such as generating separate training and test sets. Second, journal editors should be made aware of this issue and enforce the requirement that test results are generated independently from training results. Third, referees should recognize this potential error and adopt the same standards as editors. Finally, researchers in automatic algorithm configuration have long understood that testing instances must be different from training instances, so another reason why one should use automated configuration [3, 4].

## References

- [1] IBM ILOG CPLEX Optimization Studio, 2024.
- [2] Gurobi Optimization, 2024.
- [3] F. Hutter, H. H. Hoos, K. Leyton-Brown, Sequential model-based optimization for general algorithm configuration, in: C. A. Coello Coello (Ed.), Learning and Intelligent Optimization, 5th International Conference, LION 5, Vol. 6683 of LNCS, Springer, 2011, pp. 507–523.
- [4] López-Ibáñez, M., Dubois-Lacoste, J., Pérez Cáceres, L., Stützle, T., and Bittartari, M. (2016). The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58.
- [5] Y. Nagata and S. Kobayashi. A powerful genetic algorithm using edge assembly crossover for the traveling salesman problem. *INFORMS Journal on Computing*, 25(2):346–363, 2013.
- [6] Stützle, T. and Hoos, H. H. (2000). *MAX-MIN* Ant System. *Future Generation Computer Systems*, 16(8):889–914.

# Models and algorithms for the time-dependent traveling salesman problem with drone

Shuo Zhang

School of Economics and Management, Beijing Jiaotong University

e-mail: 24120556@bjtu.edu.cn

Wenchao Wei

School of Economics and Management, Beijing Jiaotong University

e-mail: weiwenchao@bjtu.edu.cn

The time-dependent traveling salesman problem with drone (TD-TSPD) is a variant of the well-known traveling salesman problem with drone (TSPD), in which travel time is a function of departure time. The TD-TSPD incorporates the time-varying congestion, and therefore is widely studied for distribution problems in large cities. The merit of teaming a truck and a drone is to take advantage of the larger capacity of the truck and the faster speed of the drone, as it is not restricted to a congested road network.

We consider a directed graph  $G = (V, A)$ . The set of vertices  $V$  is defined as  $V = \{0, 0'\} \cup N$ , where  $0$  and  $0'$  represent the same depot and  $N$  is a set of customers to be served. The set of arcs  $A$  is defined as  $A = \{(0, j) | j \in N\} \cup \{(i, 0') | i \in N\} \cup \{(i, j) | i, j \in N, i \neq j\}$ . A single truck and drone are deployed at the depot and are used to serve all customers  $N$ . The trip takes place within a planning horizon  $[0, T]$ . For each  $(v, w) \in A$ , the time required for the truck to travel from  $v$  to  $w$  is given by a continuous, piecewise linear function  $\tau_{vw}^T(t)$  [2]. Specifically, for each  $t \leq T$ ,  $\tau_{vw}^T(t)$  is the time the vehicle travels from  $v$  to  $w$  when it leaves  $v$  at time  $t$ . Note that  $\tau_{vw}^T(t)$  satisfies the first-in, first-out property (FIFO,  $\tau_{vw}^T(t) < \tau_{vw}^T(t')$  for every  $t < t'$ ). The time to traverse an arc  $(v, w)$  with the drone is a constant  $\tau_{vw}^D$  and, in general, we have  $\tau_{vw}^D \leq \tau_{vw}^T$ . Each customer can be served by the truck alone, by the drone alone, or by both the truck and the drone teamed up. Due to its limited battery capacity, the drone can serve at most one customer per trip before returning to the truck. The drone is permitted to launch from and return to the truck at any customer node or the depot, with the restriction that a take-off and subsequent landing cannot occur at the same node. The truck may wait at a customer point for the drone, whereas the drone is not allowed to wait there for the truck. The objective of TD-TSPD is to determine a coordinated tour for the truck and drone that minimizes the total completion time, where each customer is served by either vehicle or drone, considering the waiting time arising from synchronization.

To the best of our knowledge, this study presents the first framework that integrates time-dependent travel times into the TSPD. We argue that incorporating time-dependency in the TSPD not only accounts for real-world characteristics, but also enhances the synchronization between truck and drone, leading to fur-

ther optimization. Following the approaches of Adamo et al. (2024) [1], Roberti and Ruthmair (2021) [3], and Tamke and Buscher (2021) [4], we formulate the TD-TSPD as a Mixed Integer Programming model. Furthermore, we develop a branch-and-cut algorithm based on this formulation—a solution strategy previously applied to the TSPD and TDTSP individually. The proposed algorithm will be evaluated on standard benchmark instances from the literature and compared with existing exact and heuristic methods.

## References

- [1] Adamo T., Gendreau M., Ghiani G., Guerriero E., 2024. A review of recent advances in time-dependent vehicle routing. *European Journal of Operational Research*, 319(1): 1-15.
- [2] Ichoua S., Gendreau M., Potvin, J.-Y., 2003. Vehicle dispatching with time-dependent travel times. *European Journal of Operational Research*, 144(2): 379-396.
- [3] Roberti R., Ruthmair M., 2021. Exact Methods for the Traveling Salesman Problem with Drone. *Transportation Science*, 55(2): 315-335.
- [4] Tamke F., Buscher U., 2021. A branch-and-cut algorithm for the vehicle routing problem with drones. *Transportation Research Part B: Methodological*, 144: 174-203.

# A Multi-Agent Approach to Improving XAI-Narratives

Yifan He

University of Antwerp, Department of Engineering Management

e-mail: [yifan.he@uantwerpen.be](mailto:yifan.he@uantwerpen.be)

David Martens

University of Antwerp, Department of Engineering Management

e-mail: [david.martens@uantwerpen.be](mailto:david.martens@uantwerpen.be)

Using LLM to translate quantitative XAI explanations, such as SHAP, into lay-user-friendly narratives is an increasingly active area of research[1]. However, this approach is not without limitations. LLMs can introduce faithfulness-related issue when generating narratives, mainly due to their stochasticity nature and counterintuitive feature influence from the original SHAP table. This issue is particularly critical because unfaithful narratives may present misleading information that is difficult to detect, possibly undermining the reliability and trustworthiness of the explanations. Existing studies have proposed several quantitative metrics to automatically assess the quality of these narratives[2]. Building on this foundation, our goal is therefore to generate and improve XAI narratives through an automated approach.

Therefore, we develop a multi-LLM agentic system to autonomously refine XAI narratives. In this framework, the Narrator (NA) interacts with several critic agents, namely the Faithful Evaluator (FFE), Faithful Critic (FFC), and Coherence Agent (CO), each providing targeted feedback that guides the refinement through iterative processes (Figure 1). We also design different agentic system variants to compare the effects: not all variants incorporate every agent, and only the NA and the FFE are included as fixed players. To further examine model-level differences, we experiment with five LLMs to determine which performs best in this task. Extraction errors are discovered to be a major source of faithfulness issues in our experiments; therefore, we also propose an ensemble mechanism: all five LLMs serve as parallel FFEs whose outputs are aggregated via majority voting to produce a more reliable extraction, thereby reducing the likelihood of extraction-related mistakes.

Overall, Claude-4.5-Sonnet outperforms the other four LLMs in improving narrative faithfulness. Nearly all agentic system designs show clear improvements, with the exception that the two variants incorporating a coherence agent fail to improve rank accuracy. For Claude-4.5-Sonnet and ChatGPT 5, the NA+FFE variant achieves the strongest results, whereas for the remaining LLMs, adding a Faithful Critic leads to the best performance. Across all settings, systems that include a coherence agent consistently achieve the weakest results, which is expected due to its role. The ensemble versions outperform their single-LLM counterparts

across all faithfulness metrics, confirming the effectiveness of the majority-voting strategy. In summary, our results demonstrate that agentic systems provide an effective and autonomous approach for improving LLM-generated XAI narratives.

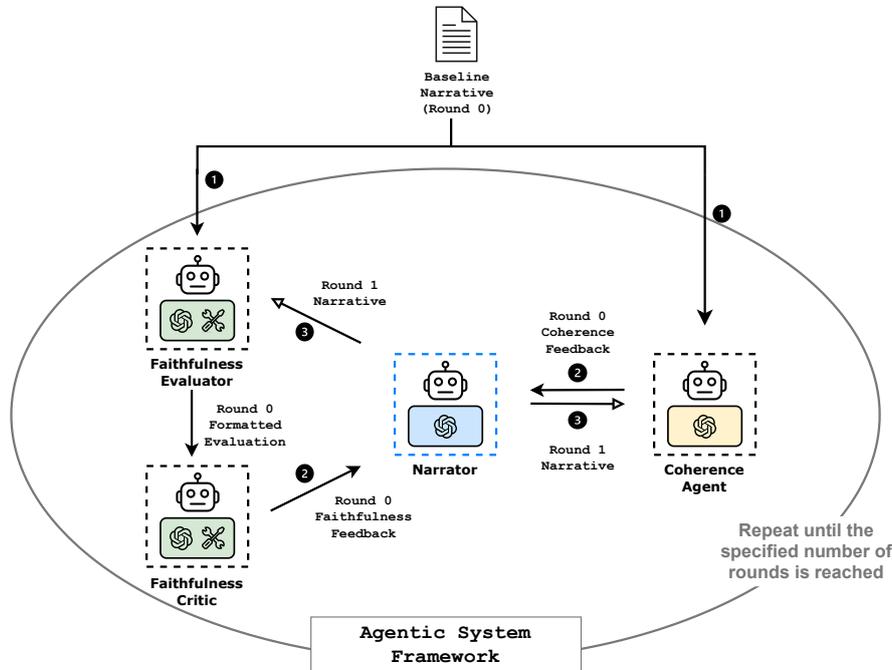


Figure 1: An overview of the agentic workflow. The baseline narrative (Round 0 narrative) is first passed to Faithful Evaluator (FFE) and Coherence Agent (CO). The FFE extracts structured information from each narrative and identifies any faithfulness issues (rank error, sign error and value error), which are then passed to the Faithful Critic (FFC). The FFC compares these issues against the SHAP table and produces targeted correction instructions. In parallel, the CO detects coherence-related problems and provides its own modification suggestions. The Narrator (NA) receives all available feedback and generates an updated narrative for the next round. This iterative process continues until the stopping criterion (three rounds) is met.

## References

- [1] D. Martens, J. Hinns, C. Dams, M. Vergouwen, T. Evgeniou, Tell Me a Story! Narrative-Driven XAI with Large Language Models, **2024**.
- [2] T. Ichmoukamedov, J. Hinns, D. Martens, How good is my story? Towards quantitative metrics for evaluating LLM-generated XAI narratives, **2024**.

# Cash or Comfort?

## How LLMs Value Your Inconvenience

Mateusz Cedro      Timour Ichmoukhamedov      Sofie Goethals  
 Yifan He              James Hinns  
 David Martens

University of Antwerp, Department of Engineering Management  
 e-mail: [mateusz.cedro@uantwerpen.be](mailto:mateusz.cedro@uantwerpen.be)

Large Language Models (LLMs) are increasingly proposed as near-autonomous artificial intelligence (AI) agents capable of making everyday decisions on behalf of humans [2, 3]. Although LLMs perform well on many technical tasks, their behaviour in personal decision-making remains less understood. Previous studies have assessed their rationality and moral alignment with human decisions. However, the behaviour of AI assistants in scenarios where financial rewards are at odds with user comfort has not yet been thoroughly explored [1].

In this paper, we tackle this problem by quantifying the prices assigned by multiple LLMs to a series of user discomforts: additional walking, waiting, hunger and pain (see Figure 1). We uncover several key concerns that strongly question the prospect of using current LLMs as decision-making assistants: (1) a large variance in responses between LLMs, (2) within a single LLM, responses show fragility to minor variations in prompt phrasing (e.g., reformulating the question in the first person can considerably alter the decision), (3) LLMs can accept unreasonably low rewards for major inconveniences (e.g., €1 to wait 10 hours), (4) LLMs can reject monetary gains where no discomfort is imposed (e.g., €1,000 to wait 0 minutes), and (5) models often make different decisions at rounded figures (e.g., €10, €100, or €1,000) (see Table 1 and Figure 2). These findings emphasise the need for scrutiny of how LLMs value human inconvenience, particularly as we move toward applications where such cash-versus-comfort trade-offs are made on users’ behalf.

Model	Time	Distance	Hunger	Pain	Avg. Value	Avg. Rank
temperature=1.0	€ for 60 min	€ for 5 km	€ for 60 min	€ for 50 %		
Gemini 2.0 Flash	0.41 $\pm$ 0.0	2.62 $\pm$ 0.3	2.26 $\pm$ 0.3	1.24 $\pm$ 0.2	1.63	1.25
Llama 3.3 70B	0.92 $\pm$ 0.1	3.41 $\pm$ 0.3	4.01 $\pm$ 0.5	1.76 $\pm$ 0.2	2.53	2.50
DeepSeek V3	2.00 $\pm$ 0.2	5.73 $\pm$ 0.5	8.71 $\pm$ 0.9	2.30 $\pm$ 0.3	4.69	3.50
Mixtral 8x22B	9.38 $\pm$ 0.8	2.86 $\pm$ 0.5	< 0.10	> 10 <sup>3</sup>	253.09	3.50
GPT-4o	5.22 $\pm$ 0.3	21.58 $\pm$ 1.7	26.36 $\pm$ 1.9	92.79 $\pm$ 15.1	36.49	5.00
Claude 3.5 Sonnet	9.76 $\pm$ 0.7	8.90 $\pm$ 0.7	50.96 $\pm$ 5.6	4.85 $\pm$ 0.5	18.62	5.25
<b>Avg. Value</b>	4.62	7.52	15.40	183.82		

Table 1: The prices of inconvenience across four scenarios for LLMs. Reported values represent the mean  $\pm$  standard deviation of the 50%-acceptance thresholds.

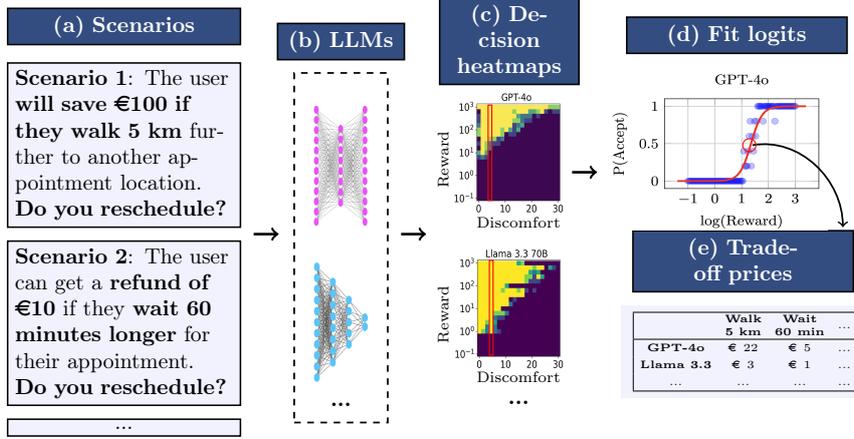


Figure 1: Methodology overview. (a) We assess four inconvenience-reward scenarios. (b) Six LLMs act as decision-making assistants for users. (c) Heatmaps display acceptance probabilities. (d) We fit a logit classifier to specific inconvenience levels (e) to derive transition points, comparing LLM valuations.

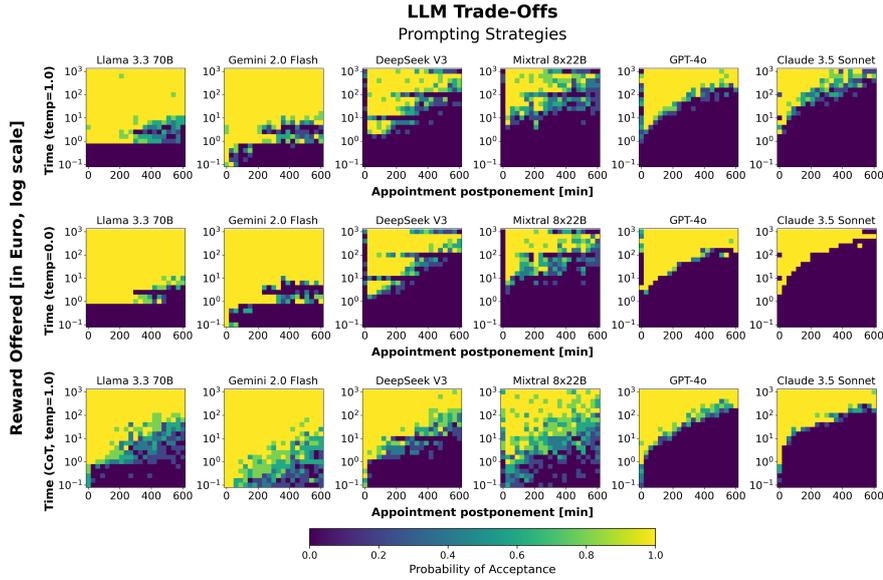


Figure 2: *Time* scenario acceptance heatmaps for six LLMs. Comparison of standard generation (Temp=1.0, 0.0) vs. Chain-of-Thought technique (Temp=1.0).

## References

- [1] Yiting Chen et al. The emergence of economic rationality of GPT. *Proceedings of the National Academy of Sciences*, 2023.
- [2] Jur Gaarlandt et al. AI agents are changing how people shop. Here’s what that means for brands. *Harvard Business Review*, 26:2, 2025.
- [3] Noam Kolt. Governing AI agents. *arXiv preprint arXiv:2501.07913*, 2025.

## Foundation Models for Credit Risk Prediction: Game Changer or False Hope?

Andreas Goethals

KU Leuven, LIRIS

e-mail: [andreas.goethals@kuleuven.be](mailto:andreas.goethals@kuleuven.be)

Bart Baesens

KU Leuven, LIRIS

e-mail: [bart.baesens@kuleuven.be](mailto:bart.baesens@kuleuven.be)

Stefan Lessmann

Humboldt-University of Berlin

e-mail: [stefan.lessmann@hu-berlin.de](mailto:stefan.lessmann@hu-berlin.de)

Simon De Vos

KU Leuven, LIRIS

e-mail: [simon.devos@kuleuven.be](mailto:simon.devos@kuleuven.be)

Victor Medina-Olivares

University of Edinburgh

e-mail: [Victor.Medina@ed.ac.uk](mailto:Victor.Medina@ed.ac.uk)

David Martens

University of Antwerp

e-mail: [david.martens@uantwerpen.be](mailto:david.martens@uantwerpen.be)

Christophe Mues

University of Southampton

e-mail: [C.Mues@soton.ac.uk](mailto:C.Mues@soton.ac.uk)

Tim Verdonck

University of Antwerp / KU Leuven

e-mail: [tim.verdonck@uantwerpen.be](mailto:tim.verdonck@uantwerpen.be)

Wouter Verbeke

KU Leuven, LIRIS

e-mail: [wouter.verbeke@kuleuven.be](mailto:wouter.verbeke@kuleuven.be)

09/12/2025

Predictive models play a crucial role in credit risk management, informing critical financial decisions through accurate estimation of default probabilities and losses given default. Their performance directly influences both the stability of the financial system and the profitability of lending operations. Large-scale benchmarking studies have periodically assessed the state of the art, systematically revealing the merits and limitations of methodological advancements. Today, tree-based ensembles such as gradient-boosting models remain the dominant approach for this application. Nevertheless, continuous improvement of these predictive models remains essential.

Concurrently, foundation models, pretrained on extensive datasets from diverse domains, have transformed predictive modelling in natural language processing and computer vision. Recently, tabular foundation models such as TabPFN have also emerged, promising strong performance without task-specific training or hyperparameter tuning. We conjecture that pretraining on out-of-domain data (synthetic data, in TabPFN's case) may prove particularly beneficial in small-data settings common in credit risk, such as SME lending or specialized corporate portfolios, and may help address longstanding challenges including low-default

portfolios and class imbalance. However, the actual value of foundation models for credit risk prediction remains an open question.

This study investigates whether these new foundation models for tabular data confer advantages in credit risk modelling. To this end, we conduct an extensive benchmark comparing TabPFN against established machine learning methods on two core tasks: Probability of Default (PD) classification and Loss Given Default (LGD) regression. Our experimental framework encompasses 15 PD datasets and 5 LGD datasets from public and proprietary sources, covering retail and small business contexts with sample sizes ranging from hundreds to over 200,000 observations. We evaluate TabPFN alongside gradient boosting machines (XGBoost, LightGBM, CatBoost), random forests, logistic regression, and deep learning baselines (MLP and TabNet). All methods are assessed using five-fold cross-validation, both with default hyperparameters and after Optuna-based tuning (50 trials per method).

Our results reveal task-specific patterns. For PD modelling, TabPFN delivers competitive performance out of the box, ranking third overall behind only tuned CatBoost and tuned Random Forest, despite receiving no hyperparameter optimization while competing methods benefited from extensive tuning. Notably, TabPFN's zero-shot performance matches or exceeds that of tuned gradient boosting methods: it outperforms tuned XGBoost and matches tuned LightGBM while requiring no hyperparameter optimization. For LGD regression, all methods struggle with the characteristic bimodal target value distributions of LGD datasets, resulting in relatively low  $R^2$  values. Yet TabPFN emerges as the top performer among tuned methods, narrowly outperforming tuned XGBoost and LightGBM while again requiring no tuning.

These findings carry practical implications: TabPFN offers a compelling option when computational resources for hyperparameter tuning are limited or when rapid model deployment is required. Unlike traditional methods, TabPFN ensures ease of use and mitigates the computational costs typically associated with model development. Our findings are relevant to academics advancing machine learning methodologies for tabular data and practitioners seeking efficient, easy-to-implement solutions for credit risk management.

# Causal Modeling of Organ Offer Dynamics for Better Allocation

Alessandro Marchese

Vrije Universiteit Brussel, Data Lab  
e-mail: [alessandro.marchese@vub.be](mailto:alessandro.marchese@vub.be)

Sam Verboven

Vrije Universiteit Brussel, Data Lab  
e-mail: [sam.verboven@vub.be](mailto:sam.verboven@vub.be)

Organ transplantation remains severely constrained by donor scarcity and inefficiencies in offer acceptance. Across organ types, acceptance rates in real-world systems commonly fall below 5%, with most offers refused multiple times before a successful placement [3]. These refusals prolong cold ischemic time (CIT), degrade graft quality, and substantially increase the probability of eventual discard [4]. Any allocation policy aiming to improve health outcomes must therefore model not only long-term transplant benefit, but also acceptance behavior, decision time, and the sequential dynamics governing organ deterioration. Yet current allocation models, including those embedded in regulatory simulators, rely on oversimplified, non-causal acceptance estimators, failing to generalize to counterfactual policies or to incorporate clinician-provided explanatory signals.

This work synthesizes two complementary advances that jointly address these structural limitations: (i) CLEXNET [1], a causal acceptance model that learns from observational data augmented with direction-only human explanations and, (ii) DYNAMITE [2], a dynamic allocation framework that optimizes match-runs by jointly modeling acceptance probabilities, decision time, CIT evolution, and individualized treatment effects (ITEs). Together, these methods form an operational pipeline for policy-relevant counterfactual estimation and optimized organ allocation.

## 1 Motivation and Background

Offer refusals are not merely noisy labels but carry meaningful, contrastive human explanations—e.g., “old donor age”, “high cold ischemic time”. These coded reasons implicitly contain counterfactual information: had the donor been younger, or the CIT lower, acceptance might have occurred. However, this information comes without magnitudes, monotonicity, or precise targets. Existing explanation-guided ML methods cannot use such weak signals. Consequently, acceptance models trained purely on observational data inherit policy-induced confounding, shortcut learning, and poor counterfactual generalization.

At the same time, allocation policies that ignore acceptance behavior—such as MELD-based or benefit-based strategies—implicitly assume deterministic, immediate acceptance. This causes systematic misestimation of ITEs for downstream candidates and ignores the compounding effects of sequential refusals on CIT. Such sequential dependencies violate the Stable Unit Treatment Value Assumption (SUTVA) and lead to biased survival estimates.

## 2 Methods

CLEXNET reformulates acceptance estimation as a counterfactual prediction problem guided by categorical refusal explanations. Each refusal code is mapped to a signed direction vector in the joint patient–donor feature space, indicating which features must increase, decrease, or remain unchanged to convert a refusal into an acceptance. Although magnitudes are unknown, CLEXNET introduces an explanation-guided augmentation loss and combines it with typical causal representation learning.

DYNAMITE introduces a policy optimization framework that explicitly models: offer acceptance probabilities, decision times, CIT evolution under sequential refusals and ITE-based transplant benefit. To support these components, DYNAMITE trains PATIENTNET, a multi-task acceptance + decision time predictor, and uses an ITE estimator to predict post-transplant benefit conditioned on dynamically updated organ quality to account for spillover effects. DYNAMITE then seeks the ranking that maximizes the expected post-transplant benefit by combining recursive organ quality updates and a local search.

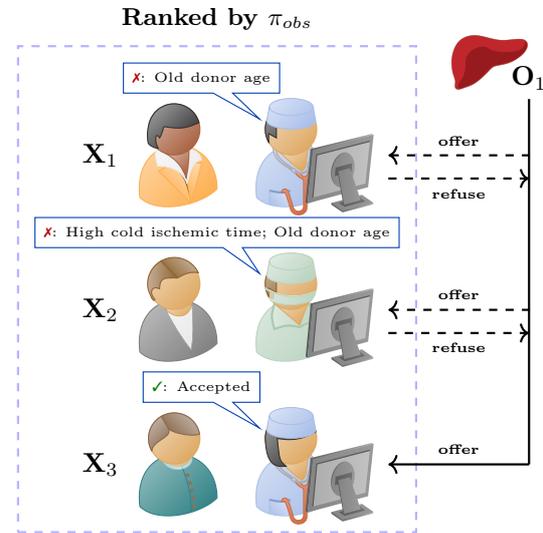


Figure 1: **Illustrative overview of organ offer acceptance.** An incoming liver offer  $O_1$  is broadcasted down the observed-policy  $\pi_{obs}$  ranking. Each candidate  $X_i$  either accepts (green tick) or refuses (red cross) and—if refusing—supplies a categorical refusal reason (call-outs).

## References

- [1] Alessandro Marchese, Jeroen Berrevoets, and Sam Verboven. Causal explanation-guided learning for organ allocation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [2] Alessandro Marchese, Hans de Ferrante, Jeroen Berrevoets, and Sam Verboven. Dynamite: Optimal time-sensitive organ offers using ite. *Machine Learning for Health (ML4H)*, pages 696–713, 2025.
- [3] Jonathan M Miller, Yoon Son Ahn, Allyson Hart, Dorry L Segev, David P Schladt, Kathryn T Livelli, Kelsi A Lindblad, Ajay K Israni, and Jon J Snyder. Optn/srtr 2022 annual data report: Covid-19. *American Journal of Transplantation*, 24(2):S489–S533, 2024.
- [4] Sumit Mohan, Mariana C. Chiles, Rachel E. Patzer, Stephen O. Pastan, S. Ali Husain, Dustin J. Carpenter, Geoffrey K. Dube, R. John Crew, Lloyd E. Ratner, and David J. Cohen. Factors leading to the discard of deceased donor kidneys in the United States. *Kidney International*, 94(1):187–198, July 2018. Publisher: Elsevier BV.

## Hierarchical forecasting and decision-making: some perspectives from the energy sector

Talia Qaiser

KU Leuven, Dept. of Electrical Engineering

e-mail: [talia.qaiser@kuleuven.be](mailto:talia.qaiser@kuleuven.be)

Fahad Mehmood

EM Normandie,

e-mail: [fmehmood@em-normandie.fr](mailto:fmehmood@em-normandie.fr)

Geert Deconinck

KU Leuven, Dept. of Electrical Engineering

e-mail: [geert.deconinck@kuleuven.be](mailto:geert.deconinck@kuleuven.be)

Hussain Kazmi

KU Leuven, Dept. of Electrical Engineering

e-mail: [hussain.kazmi@kuleuven.be](mailto:hussain.kazmi@kuleuven.be)

Electricity grid operators are facing unprecedented challenges in safe and secure power system operation. At the distribution grid level, the rapid proliferation of distributed generation (e.g. rooftop solar PV systems) can lead to large reverse power flows. Likewise, accelerating demand electrification (e.g. heating and transportation) can further exacerbate grid congestion events and voltage issues [1].

Therefore, relying solely on the historical fit-and-forget paradigm (i.e. investing in grid reinforcement alone) will no longer suffice for (distribution) grid operators mandated to ensure a certain quality of service. Instead, they must now increasingly rely on operationalizing customer demand flexibility to alleviate the worst effects of distributed generation and demand electrification [2].

To this end, across Europe, the distribution system operators (DSOs) are investing heavily in activating end-user flexibility. A crucial first step is accurately quantifying and forecasting base energy demand—the driver of flexibility. This can be done at the customer level, where flexibility is activated, or at the aggregated feeder level, where DSOs typically have measured data. [3].

Operating at these two levels in isolation can, however, lead to incoherent forecasts and decisions: the DSO's forecast (the forecast of total demand of all customers) is not necessarily the same as the customer-side forecasts (the sum of individual customer demand forecasts). In our experiments with real-world smart meter data, we have seen this coherency gap to consistently range between 5 and 10%, even ignoring all other technical and non-technical losses. Such a gap can easily mean the difference between a correct and incorrect decision to (not) activate flexibility.

Hierarchical forecasting has been proposed as a potential solution to this problem, it reconciles the forecasts made at various levels [4]. For instance, MinT, a popular technique, reconciles these forecasts in a way that minimizes the variance of the reconciled forecast errors [5]. Often, the reconciliation process also results in a (small) boost to predictive accuracy, as the global information available at top hierarchical levels allows for better calibration of individual, lower level forecasts (which tend to be noisier).

The hierarchical forecasting problem can be written as  $\tilde{y}_h = SG\hat{y}_h$ , where  $\hat{y}_h$  represents the base unreconciled forecasts,  $h$  is the forecast horizon,  $G$  is a matrix mapping the base forecasts into the bottom level, and  $S$  is a matrix that sums the unreconciled forecasts using the aggregation structure to produce a set of coherent forecasts  $\tilde{y}_h$ .

However, despite considerable research in hierarchical forecasting, current methods do not represent a panacea for DSOs. One reason for this is that distribution grids, especially when considering customers connected to a single feeder, represent very shallow, single-level hierarchies. In such cases, we hypothesize that the reconciliation process does not bring too much benefit because of the limited cross-scale information flow and difficulties in estimating the  $h$ -step ahead base forecast error variance.

To address this problem, we propose a unified method that integrates hierarchical clustering and forecasting by first extracting a dendrogram from the base level time series, and uses this to construct the summing matrix  $S$  in a data-driven way. The introduction of several new hierarchical levels enables the reconciliation process to result in coherent, more accurate forecasts despite the originally shallow nature of the hierarchy.

Preliminary results, using a large-scale electricity demand dataset from London, show that the proposed method significantly outperforms reconciliation strategies that do not group time series together a priori. Averaged-out coherency gap without reconciliation was 7.1%, after hierarchical clustering and MinT(OLS) reconciliation the gap closes to 3.3% and MinT(WLS) to about 3.4%. These results also shed fresh light on how hierarchical approaches can be linked to global forecasting methods, which model multiple time series jointly, as well as their implications on downstream decision-making.

## References

- [1] C. Protopapadaki and D. Saelens, “Heat pump and PV impact on residential low-voltage distribution grids as a function of building and district properties,” *Applied Energy*, vol. 192, pp. 268–281, 2017.
- [2] J. Le Dréau, R. A. Lopes, S. O’Connell, D. Finn, M. Hu, H. Queiroz, D. Alexander, A. Satchwell, D. Österreicher, B. Polly *et al.*, “Developing energy flexibility in clusters of buildings: A critical analysis of barriers from planning to operation,” *Energy and Buildings*, vol. 300, p. 113608, 2023.
- [3] H. Kazmi, C. Fu, and C. Miller, “Ten questions concerning data-driven modelling and forecasting of operational energy demand at building and urban scale,” *Building and Environment*, vol. 239, p. 110407, 2023.
- [4] G. Athanasopoulos, R. J. Hyndman, N. Kourentzes, and A. Panagiotelis, “Forecast reconciliation: A review,” *International Journal of Forecasting*, vol. 40, no. 2, pp. 430–456, 2024.
- [5] S. L. Wickramasuriya, G. Athanasopoulos, and R. J. Hyndman, “Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization,” *Journal of the American Statistical Association*, vol. 114, no. 526, pp. 804–819, 2019.

# Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand

Sarah Van der Auweraer

IESEG School of Management, e-mail: [s.vanderauweraer@ieseg.fr](mailto:s.vanderauweraer@ieseg.fr)

Thomas van Pelt

University of Luxembourg, e-mail: [tdvpelt@outlook.com](mailto:tdvpelt@outlook.com)

Joachim Arts

University of Luxembourg, e-mail: [joachim.arts@uni.lu](mailto:joachim.arts@uni.lu)

4 December 2025

Intermittent demand is characterized by many periods with no demand at all, that are seemingly randomly interspersed with demand occurrences. It is unknown when the next demand will occur, and there is uncertainty about the size of demand when it occurs. This compound nature of demand uncertainty complicates efficient inventory management.

Specific approaches have been developed to deal with such intermittent demand. In the existing literature, the analysis of intermittent demand times series is commonly decomposed to create separate estimates for the time between demand occurrences (the inter-demand interval) and the size of a demand occurrence, initiated by the work of Croston [1]. These analyses often implicitly assume that the time between demand occurrences is Markovian, i.e., the probability of observing a demand in a certain period is considered to be independent of the time since the last demand occurrence. However, data from practice indicate that the times between demand occurrences are not necessarily Markovian. Non-Markovian times between demand occurrences lead to *rhythmic* patterns in demand processes when the hazard rate is increasing, which has recently also been observed by [2] and [3]. Consequently, the time since the last demand occurrence can be an important predictor for future demand for those items.

In this work we propose a demand model that accommodates such rhythmic intermittent demand. We model demand as a discrete compound renewal process. Unsatisfied demand is backordered at a penalty cost and inventory on-hand at the end of a period incurs a holding cost. We allow for positive leadtimes.

This research is strongly rooted in practice. It is the outcome of a joint research project with a company in the chemical industry, which produces specialty chemicals for its customers within a make-to-stock environment. The company experiences rhythmic intermittent demand for many of its items, as observed in their data. Nonetheless, the company currently does not consider this rhythm in its demand forecasting or inventory decision making.

The model we developed in the context of this collaboration contributes to

both the theory and practice of inventory management as we characterize the structure of optimal replenishment policies and design heuristics based on these structural results. The main contributions of our work include descriptions of the form of optimal order policies. In particular, we show that the optimal policy is a state-dependent base-stock policy, where the state is the time since the last demand observation. We prove this by induction on the value function. We also show that there exist state-dependent base-stock policies for which the optimal base-stock level is non-decreasing in the time since the last demand. As such, any algorithm that searches for optimal (or good) base-stock levels can constrain the search space to non-decreasing base-stock levels only. Contrary to what may be expected, this result is not proven by showing submodularity of the value function. Instead we prove sensitivities of optimal base-stock levels with elementary arguments, which can provide new perspectives. We characterize the sensitivities of the base-stock level to the current state and the proof exploits structure in compound renewal processes directly.

Moreover, we also contribute to inventory control practice. We demonstrate how the proposed approach can be implemented in practice through the development of numerically stable algorithms that can scale to large assortments as encountered in industry. Additionally, we introduce two heuristic policies that are easier to implement

We benchmark the performance of the different approaches both in a numerical experiment in which the demand process is assumed to be known, and on real data, where the true demand process is not known but demand data is available. We find that using information on the demand rhythm and the time since the last demand occurrence is valuable when demand occurs at regular intervals, with limited deviations. Additionally, when demand is more intermittent, using such information is increasingly beneficial. Nonetheless, a trade-off with data availability exists.

## References

- [1] J.D. Croston, 1972, Forecasting and Stock Control for Intermittent Demands, *Operational Research Quarterly*, 23(3): 289–303.
- [2] D. Prak and P. Rogetzer, 2022, Timing intermittent demand with time-varying order-up-to levels, *European Journal of Operational Research*, 303(3): 1126–1136.
- [3] A.C. Turkmen, T. Januschowski, Y. Wang, and A.T. Cemgil, Intermittent Demand Forecasting with Renewal Processes, arXiv preprint arXiv:2010.01550.

## On the impact of forecast reliability on downstream decision-making

Jente Van Belle

KU Leuven, Faculty of Economics and Business,  
Department of Decision Sciences and Information Management  
e-mail: [jente.vanbelle@kuleuven.be](mailto:jente.vanbelle@kuleuven.be)

Wouter Verbeke

KU Leuven, Faculty of Economics and Business,  
Department of Decision Sciences and Information Management  
e-mail: [wouter.verbeke@kuleuven.be](mailto:wouter.verbeke@kuleuven.be)

Nikolaos Kourentzes

Independent

e-mail: [nikolaos@kourentzes.com](mailto:nikolaos@kourentzes.com)

While assessing forecast accuracy is of key importance in the evaluation of time series forecasting models, the consistency of their performance over time has remained largely underexplored in the literature. We refer to this property as temporal forecast reliability.

A common evaluation approach is the rolling-origin setup, in which forecast traces containing one- to  $h$ -step-ahead forecasts are generated from multiple forecasting origins [1]. By aggregating performance metrics across these origins, the evaluation captures how methods adapt to varying conditions and accounts for multiple test periods [2]. However, this is typically limited to mean performance across forecast traces and thus overlooks the consistency of performance over time: occasional unreliable forecasts may be masked by these summary statistics. For example, an analysis of performance over time may reveal that one method has stationary and homoscedastic errors, while another exhibits periods of highly volatile or nonstationary performance. Avoiding large errors is crucial, as these can substantially impact downstream operations and erode confidence in the forecasting system [3]. However, identifying which forecasting methods generate reliable predictions—and which do not—requires moving beyond merely assessing mean performance over time [4].

In this work, we formalize the concept of temporal forecast reliability and demonstrate its importance as an additional evaluation criterion. Specifically, we show that when forecasts are used as inputs to decision-making processes, differences in forecast reliability can affect downstream decision-making performance due to a nonlinear propagation of forecast errors. As operational decision-making becomes increasingly automated, we argue that forecast reliability should be treated as a key evaluation dimension rather than an afterthought.

## References

- [1] Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: An analysis and review. *International Journal of Forecasting*, 16(4), 437–450.
- [2] Athanasopoulos, G., & Kourentzes, N. (2023). On the evaluation of hierarchical forecasts. *International Journal of Forecasting*, 39(4), 1502–1511.
- [3] Spavound, S., & Kourentzes, N. (2022). Making forecasts more trustworthy. *Foresight: The International Journal of Applied Forecasting*, 66, 21–25.
- [4] Seedat, N., Imrie F., & van der Schaar, M. (2024). Navigating data-centric artificial intelligence with DC-check: Advances, challenges, and opportunities. *IEEE Transactions on Artificial Intelligence*, 5(6), 2589–2603.

# Semantic information for variable selection in tactical business forecasting

Yves R. Sagaert

VIVES University of Applied Sciences, Computer Science

e-mail: [yves.sagaert@vives.be](mailto:yves.sagaert@vives.be)

Nikolaos Kourentzes

e-mail: [nikolaos@kourentzes.com](mailto:nikolaos@kourentzes.com)

Leading indicators have been shown to be useful in predicting demand. Nowadays, a vast number of potentially informative covariates are accessible through open data repositories, which substantially complicates the identification of explainable predictors with predictive value. This difficulty increases in complexity when variables exhibit strong multicollinearity or when their dimensionality vastly exceeds the available estimation sample size, conditions under which variable selection quality deteriorates. Both phenomena are highly pertinent in practice: in aggregate-demand modelling, the influence of exogenous variables typically becomes more pronounced, while the space of potential predictors expands considerably [4]. Lasso regression has been widely adopted for handling high-dimensional covariate spaces. However, its variable selection is fragile with respect to sampling variation and is particularly sensitive in the presence of correlated predictors. This instability can undermine both predictive accuracy and the perceived credibility and reliability of the resulting forecasts.

Here, we want to address these limitations by structuring the input variables into semantically coherent groups. This approach builds upon two strands of prior literature: first, empirical evidence that lasso-based demand forecasts yield strong performance in tactical and strategic planning contexts [2, 3]; and second, findings that grouping variables using domain expertise or prior information can enhance both variable selection and forecast accuracy [1, 2, 6]. Several strategies are investigated to support variable selection: grouping based on statistically driven clustering procedures, semantic representations derived from a Semantic Bidirectional Encoder Representations from Transformers (BERT) model, and meta-data features such as a popularity index of variables. The resulting variable groups are then exploited in multiple ways: (i) used directly as candidates in a sequential lasso procedure, (ii) transformed into cluster profiles, or (iii) reduced to latent factors via principal component analysis. Using an empirical case study, we compare these alternatives in terms of decision-oriented performance, predictive accuracy, and model interpretability.

While semantic information is found to be beneficial for variable selection, the analysis reveals inherent trade-offs among predictive performance, decision quality, and interpretability. Factor-based approaches deliver the best aggregate demand forecasts for tactical planning; however, they exhibit limited inter-

pretability and depend predominantly on empirical validation, akin to traditional univariate models. Moreover, their forecasting performance deteriorates when applied to more disaggregated demand series, which constrains their generalisability. In contrast, approaches that integrate semantic information yield more consistently positive effects and enhance model explainability. In particular, the use of SBERT for variable selection appears especially promising and points to new research avenues on embedding semantic representations, encompassing both alternative uses of SBERT and the incorporation of more recent large language models (LLM), within forecasting model architectures. While analyst-generated meta-data (e.g., variable popularity) also contributes to improved performance, its impact is comparatively weaker.

Forecasts only create business value if users trust them. Otherwise, they risk of being overridden, regardless of accuracy. Trustworthiness can be expressed in four key dimensions: (1) reliability (few large errors and clear error sources), (2) stability (similar behaviour across comparable cases), (3) intelligibility (explainable key drivers), and (4) alignment with user objectives [5]. Methods with limited explainability often perform poorly on the first three, making their adoption hard to justify empirically. The integration of semantically enriched variable selection with explainable forecasting models can substantially enhance the transparency of scenario analyses in business planning. Such models should match or outperform non-explainable benchmarks, but this depends on rigorous explainability methodologies, which are still underdeveloped.

## References

- [1] Ma, S., Fildes, R., Huang, T., 2016. Demand forecasting with high dimensional data: The case of sku retail sales forecasting with intra-and inter-category promotional information. *European Journal of Operational Research* 249, 245–257.
- [2] Sagaert, Y.R., Aghezzaf, E.H., Kourentzes, N., Desmet, B., 2018. Tactical sales forecasting using a very large set of macroeconomic indicators. *European Journal of Operational Research* 264, 558–569.
- [3] Sagaert, Y.R., Kourentzes, N., De Vuyst, S., Aghezzaf, E.H., Desmet, B., 2019. Incorporating macroeconomic leading indicators in tactical capacity planning. *International Journal of Production Economics* 209, 12–19.
- [4] Sagaert, Y.R., Kourentzes, N., 2025. Inventory management with leading indicator augmented hierarchical forecasts. *Omega* 136, 103335.
- [5] Spavound, S., Kourentzes, N., 2022. Making forecasts more trustworthy. *Fore-sight: The International Journal of Applied Forecasting*, 21–25.
- [6] Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67, 301–320.

# Unified Multi-Task Predictive Process Monitoring with Transformers and Adaptive Loss Optimization

Brecht Wuyts

KU Leuven, Research Center for Information Systems Engineering (LIRIS)

e-mail: [brecht.wuyts@kuleuven.be](mailto:brecht.wuyts@kuleuven.be)

Seppe vanden Broucke

Ghent University, Department of Business Informatics and Operations Management

e-mail: [seppe.vandenbroucke@ugent.be](mailto:seppe.vandenbroucke@ugent.be)

Jochen De Weerd

KU Leuven, Research Center for Information Systems Engineering (LIRIS)

e-mail: [jochen.deweerd@kuleuven.be](mailto:jochen.deweerd@kuleuven.be)

Predictive Process Monitoring (PPM) leverages historical event logs to anticipate the future evolution of running process instances, enabling organizations to intervene before delays, failures, or undesired outcomes materialize. While contemporary deep learning models have achieved strong results across all PPM prediction tasks, most existing approaches either focus on a single task in isolation or adopt simplistic multi-task learning formulations with static, equally weighted loss functions. As a result, they remain susceptible to loss-scale disparities, heterogeneous label noise, and suboptimal optimization dynamics. At the same time, no prior work has unified all primary PPM prediction targets within a single architecture.

This work addresses these limitations by introducing *SuTraN+*, a unified Transformer-based encoder–decoder architecture that extends *SuTraN* [1] to jointly predict the full activity suffix, the corresponding timestamp suffix, the remaining runtime, and the case outcome. Beyond the architectural contribution, the paper provides the first systematic investigation of multi-task optimization (MTO) strategies in PPM, moving beyond equal weighting.

We benchmark naive equal loss weighting against several adaptive MTO techniques drawn from the broader multi-task learning literature, including uncertainty-based loss weighting [2], GradNorm [3], and PCGrad [4]. All strategies are evaluated under a strictly controlled experimental design, using identical *SuTraN+* backbones, preprocessing pipelines, and training budgets, and are compared against strong single-task baselines, leveraging the same architectural backbone as well. Experiments are conducted on four large-scale real-world event logs from heterogeneous domains.

The results demonstrate that multi-task learning consistently improves overall predictive performance compared to single-task training, even when using simple

uniform loss weighting, provided that appropriate normalization and preprocessing are applied. Among the adaptive strategies, uncertainty-based weighting emerges as the most robust and consistently effective approach across datasets. In contrast, gradient-conflict mitigation techniques offer limited benefits, suggesting that destructive gradient interference is not the dominant challenge in multi-task PPM.

A detailed analysis of learned task weights, gradient behavior, and internal optimization signals reveals that heterogeneous homoscedastic noise levels across tasks, rather than gradient conflicts, constitute the primary bottleneck in joint process prediction. These findings yield actionable guidance for practitioners, highlighting uncertainty-aware optimization as a principled and practical default for multi-task PPM.

## References

- [1] Wuyts, B., Vanden Broucke, S., & De Weerd, J. (2024, October). SuTraN: an Encoder-Decoder Transformer for Full-Context-Aware Suffix Prediction of Business Processes. In 2024 6th International Conference on Process Mining (ICPM) (pp. 17-24). IEEE.
- [2] A. Kendall, Y. Gal, and R. Cipolla, “Multi-task learning using uncertainty to weigh losses for scene geometry and semantics,” in Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 7482–7491.
- [3] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, “Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks,” in International conference on machine learning. PMLR, 2018, pp. 794–803.
- [4] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, “Gradient surgery for multi-task learning,” *Advances in neural information processing systems*, vol. 33, pp. 5824–5836, 2020.

# Causal Estimation of Midday Stem Water Potential for Decision-Oriented Irrigation Scheduling

Joaquín Roa

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [joaquin.roallanos@kuleuven.be](mailto:joaquin.roallanos@kuleuven.be)

Jan Diels

KU Leuven, Soil and Water Management

e-mail: [jan.diels@kuleuven.be](mailto:jan.diels@kuleuven.be)

Estefanía Serral A.

KU Leuven, Research Centre for Information Systems Engineering (LIRIS)

e-mail: [estefania.serralasensio@kuleuven.be](mailto:estefania.serralasensio@kuleuven.be)

December 3rd, 2025

Stem water potential (SWP) is regarded as the most reliable physiological indicator of tree water stress<sup>1</sup>. In practice, irrigation decisions aim to maintain midday SWP above a critical threshold beyond which water stress becomes significant. However, the sparse and costly nature of SWP measurements<sup>2</sup> poses substantial challenges for data-driven irrigation scheduling.

In this study, we investigate how to estimate the causal effect of irrigation and environmental covariates on SWP using high-resolution records collected in an experimental framework environment. We develop a mathematical, causally grounded formulation for predicting midday SWP, capturing how environmental phenomena, irrigation events, and soil water processes jointly determine the evolution of tree water status.

We formalize the problem as a dynamic decision setting under partial observability: at a given time, the decision-maker must anticipate SWP at a later horizon using only covariates available up to that moment. To avoid post-treatment leakage and ensure credible causal interpretation, we construct time-indexed regression models that restrict the information set to pre-decision variables and explicitly encode the temporal ordering between irrigation, soil water potential, environmental drivers, and tree responses, which leads to a system of structural

---

<sup>1</sup>Van Beek, J., Tits, L., Somers, B., & Coppin, P. (2013). Stem water potential monitoring in pear orchards through WorldView-2 multispectral imagery. *Remote Sensing*, 5(12), 6647-6666.

<sup>2</sup>Deforce, B., Baesens, B., Diels, J., & Serral, E. (2023, November). MultiMix TFT: a multi-task mixed-frequency framework with temporal fusion transformers. In *Conference on Lifelong Learning Agents* (pp. 586-600). PMLR.

components that reflect the underlying causal mechanisms that irrigation influences.

Preliminary results show a trade-off between operational feasibility and predictive accuracy. The proposed models capture key aspects of tree water uptake dynamics, including how prior irrigation and environmental conditions shape mid-day SWP. We quantitatively characterize the distinction between models that rely on earlier decision windows, which use only the information realistically available to practitioners, but produce lower predictive accuracy, and models that incorporate data closer to the measurement time, which achieve higher accuracy but leave less time margin for manual decision making. By making this trade-off explicit, our work contributes to the development of interpretable decision-support systems for irrigation scheduling that balance practical constraints with predictive performance under uncertainty.

# SCOPE: Sequential Causal Optimization of Process Interventions

Jakob De Moor

LIRIS, KU Leuven, Belgium

Hans Weytjens

LIRIS, KU Leuven, Belgium

School of Computation, Information and Technology, TUM, Germany

Johannes De Smedt

LIRIS, KU Leuven, Belgium

Jochen De Weerd

LIRIS, KU Leuven, Belgium

PresPM uses machine learning to provide case-specific recommendations at different decision points during the execution of business processes. These recommendations concern interventions, such as managerial escalations or customer communications, that aim to improve KPIs, for example, throughput time or cost efficiency. PresPM holds the potential to move organizations from merely describing and predicting process behavior towards actively steering process executions.

In many processes, intervention decisions are not isolated. Typically, a process contains multiple, interdependent decision points that jointly determine the outcome of a case. The effect of an earlier intervention depends on which interventions will be applied later in the same case. Optimizing intervention decisions one by one is therefore insufficient, because actions that look beneficial locally can undermine overall KPI performance. For example, in a marketing process aimed at maximizing revenue, the decision to offer a client a discount may depend on an earlier choice to send a promotional email and the client's response to that email. This sequence of decisions together shapes the revenue. Optimizing each decision independently without considering the future might make it look like sending a promotional email is only moderately valuable for some clients, even though if you follow it with a discount, it could be highly effective. Methods for PresPM thus need to reason over sequences of decisions and their combined effect on the final KPI.

Most existing PresPM approaches do not yet offer such sequential support. First, many methods address only a single intervention scenario, even when multiple opportunities to intervene exist [2, 5]. These approaches may improve performance for a single intervention scenario, but they do not coordinate multiple decisions over the full case. Second, some approaches handle sequential decisions but optimize each decision point in isolation. They focus on the immediate effect

of the next intervention, for instance, on the remaining processing time, without explicitly aligning interventions across decision points. This can still lead to suboptimal end-to-end outcomes with respect to the final KPI [4]. Third, other methods for sequential intervention recommendations rely on process approximations, such as Markov decision processes (MDP) or data augmentation [1, 3], to train an RL agent, and the resulting policies inherit any misspecification in these approximations, which may lead to biased or underperforming recommendations in practice. To address these limitations, we make two key contributions:

1. We propose *SCOPE*, a PresPM approach that combines causal learners with backward induction to learn causally grounded sequential intervention policies that are aligned across multiple decision points. For each decision point, a causal model estimates the effect of alternative interventions on the target KPI, given the observed process execution history. Backward induction then propagates the impact of later intervention decisions back to earlier ones by starting from the final decision point and recursively deriving the recommended action and its expected outcome at each preceding decision point. Operating directly on observational event logs, *SCOPE* uses causal learning to be able to identify actions rarely chosen historically but likely to enhance the target KPI, without requiring process-specific simulators or log augmentation.
2. We provide an empirical evaluation on existing synthetic data and a new semi-synthetic setup based on a real-life event log. Our code and the novel semi-synthetic benchmark are made publicly available as a reusable resource for sequential PresPM in our GitHub repository<sup>1</sup>. The results demonstrate that *SCOPE* consistently outperforms state-of-the-art PresPM techniques and highlight the importance of combining causal learning with backward induction for KPI optimization.

## References

- [1] M. Abbasi, M. Khadivi, M. Ahang, P. Lasserre, Y. Lucet, H. Najjaran, 2025, An Innovative Data-Driven and Adaptive Reinforcement Learning Approach for Context-Aware Prescriptive Process Monitoring.
- [2] Z. D. Bozorgi, M. Dumas, M. La Rosa, A. Polyvyanyy, M. Shoush, I. Teinmaa, 2023, Learning When to Treat Business Processes: Prescriptive Process Monitoring with Causal Inference and Reinforcement Learning.
- [3] S. Branchi, A. Buliga, C. Di Francescomarino, C. Ghidini, F. Meneghello, M. Ronzani, 2023, Recommending the optimal policy by learning to act from temporal data.
- [4] M. de Leoni, M. Dees, L. Reulink, 2020, Design and Evaluation of a Process-aware Recommender System based on Prescriptive Analytics, 2nd Int. Conf. Process Min..
- [5] M. Shoush, M. Dumas, 2025, White box specification of intervention policies for prescriptive process monitoring, Data Knowl. Eng.

---

<sup>1</sup><https://github.com/JakobDeMoorKULstudent/SCOPE>

# Handling Multi-Class Imbalance in Next Activity Prediction for Predictive Process Monitoring

Xiaomeng He

KU Leuven, LIRIS

e-mail: [xiaomeng.he@kuleuven.be](mailto:xiaomeng.he@kuleuven.be)

Rafael Oyamada

KU Leuven, LIRIS

e-mail: [rafael.oyamada@kuleuven.be](mailto:rafael.oyamada@kuleuven.be)

Johannes De Smedt

KU Leuven, LIRIS

e-mail: [johannes.desmedt@kuleuven.be](mailto:johannes.desmedt@kuleuven.be)

Seppe vanden Broucke

Ghent University, Department of Business Informatics and Operations Management

KU Leuven, LIRIS

e-mail: [seppe.vandenbroucke@ugent.be](mailto:seppe.vandenbroucke@ugent.be)

Jochen De Weerd

KU Leuven, LIRIS

e-mail: [jochen.deweerd@kuleuven.be](mailto:jochen.deweerd@kuleuven.be)

The widespread use of information systems has enabled business processes to generate extensive digital footprints, providing rich data for process mining, a data-driven approach to analyzing and optimizing business processes. Predictive Process Monitoring (PPM) is a subfield of process mining that focuses on predicting the future behavior of ongoing process instances [1]. PPM models typically rely on core control-flow attributes including case ID, activity label, and timestamp, while additional event-level or case-level attributes enable data-aware predictive modeling, which has attracted growing research interest [2].

Among the various PPM tasks, next activity prediction involves predicting the activity label of the subsequent event given an incomplete trace. Such predictions can support resource planning and enable timely interventions based on anticipated actions. Recent work has increasingly adopted machine learning (ML) and deep learning (DL) models for this task [3]. These models formulate next activity prediction as a multi-class classification problem, where each class represents a possible next activity label.

In many real-world processes, some activities tend to occur more frequently than others, which causes their corresponding classes to dominate the data and creates multi-class imbalance. Under this imbalance, standard ML models tend to be biased toward frequent classes while underperforming on infrequent ones [4]. Consequently, existing methods often struggle to deliver reliable predictions for rare but potentially critical process behavior, such as unusual cancellations or rejections. Despite its practical relevance, multi-class imbalance has received limited attention in the PPM literature. To the best of our knowledge, only

Mehdiyev et al. [5] propose a dedicated solution; however, their binary classification approach focuses on a single rare activity, limiting its applicability to multi-class settings. Appropriate approaches to address multi-class imbalance in next activity prediction are still lacking in the PPM literature.

To address this gap, this study investigates several approaches for mitigating multi-class imbalance in next activity prediction. First, motivated by the strong performance of xLSTMs in rare token prediction [6], we assess whether certain model architectures are inherently more robust to imbalance by benchmarking XGBoost, LSTM, Transformer, and xLSTM models. Second, we compare data-aware PPM models with control-flow-only models to examine whether additional event or case attributes improve the prediction of infrequent activities. Finally, we evaluate the effectiveness of loss re-weighting [7, 8] as a strategy to mitigate bias induced by class imbalance.

## References

- [1] Ceravolo, P., Comuzzi, M., De Weerd, J., Di Francescomarino, C., Maggi, F. M. (2024). Predictive process monitoring: concepts, challenges, and future research directions. *Process Science*, 1(1), 2.
- [2] Gunnarsson, B. R., vanden Broucke, S., De Weerd, J. (2023). A direct data aware LSTM neural network architecture for complete remaining trace and runtime prediction. *IEEE Transactions on Services Computing*, 16(4), 2330-2342.
- [3] Rama-Maneiro, E., Vidal, J. C., Lama, M. (2021). Deep learning for predictive business process monitoring: Review and benchmark. *IEEE Transactions on Services Computing*, 16(1), 739-756.
- [4] Johnson, J. M., Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of big data*, 6(1), 1-54.
- [5] Mehdiyev, N., Fettke, P., Evermann, J. (2020). A Novel Business Process Prediction Model Using a Deep Learning Method. *Business & information systems engineering*, 62(2), 143-157.
- [6] Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., Hochreiter, S. (2024). xlstm: Extended long short-term memory. *Advances in Neural Information Processing Systems*, 37, 107547-107603.
- [7] Cui, Y., Jia, M., Lin, T. Y., Song, Y., Belongie, S. (2019). Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9268-9277).
- [8] Lin, T. Y., Goyal, P., Girshick, R., He, K., Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision* (pp. 2980-2988).

## Parallel Machine Scheduling with Speed Scaling

Kang Yang

National University of Defense Technology, China, College of Systems Engineering

KU Leuven, Belgium, ORSTAT, Faculty of Economics and Business

e-mail: [youngclover23@outlook.com](mailto:youngclover23@outlook.com)

Roel Leus

KU Leuven, Belgium, ORSTAT, Faculty of Economics and Business

e-mail: [roel.leus@kuleuven.be](mailto:roel.leus@kuleuven.be)

In classical parallel machine scheduling problems, the processing times of jobs are typically assumed to be fixed. In this paper, we consider a parallel machine scheduling problem with adjustable processing speeds, aiming to minimize the energy consumption cost during machine operation. We refer to this problem as parallel machine scheduling with speed scaling (PMSSC) and formulate it as a convex mixed-integer nonlinear programming (MINLP) model. Leveraging the convexity of the objective function and the presence of semi-continuous variables, we strengthen the formulation using the perspective reformulation technique and develop a perspective Benders decomposition algorithm to address its computational challenges. Furthermore, we study the special case of identical parallel machines and employ symmetry-breaking techniques to enhance computational performance.

# A Decomposition Algorithm For the Integrated Flexible Job Shop Scheduling and Operator Timetabling

Reza Ghorbani Saber

Ghent University, Industrial Systems Engineering and Product Design

e-mail: [Reza.GhorbaniSaber@Ugent.be](mailto:Reza.GhorbaniSaber@Ugent.be)

Pieter Leyman

Ghent University, Industrial Systems Engineering and Product Design

e-mail: [Pieter.Leyman@Ugent.be](mailto:Pieter.Leyman@Ugent.be)

El-Houssaine Aghezzaf

Ghent University, Industrial Systems Engineering and Product Design

e-mail: [Elhoussaine.Aghezzaf@Ugent.be](mailto:Elhoussaine.Aghezzaf@Ugent.be)

This study addresses an integrated Flexible Job Shop–Operator Timetabling (FJS–OT) problem that simultaneously considers multi-shift operations, limited operators, operators–machines competencies, and machine–shift operator assignments. The Flexible Job Shop Scheduling Problem (FJSP) is a practical machine scheduling optimization problem which has received increasing attention due to its strong compatibility with modern automated manufacturing systems, such as flexible manufacturing systems. However, in real-world production environments, the feasibility of machine schedules strongly depends on operator timetabling regulations, which restrict machine availability across shifts. Integration of operator timetabling in manufacturing environments has received limited attention in previous research, especially with realistic restrictive assumptions, resulting in scheduling solutions that may be infeasible from a workforce management perspective [Artigues, Gendreau, Rousseau, and Vergnaud, 2009, Guyon, Lemaire, Pinson, and Rivreau, 2014]. The problem studied in this research work, incorporates a combination of several practical characteristics that distinguish it from existing works and reflect essential constraints observed in industrial settings:

- The production system follows a flexible job shop structure, which is more general and practical than the classical job shop setting.
- Once an operation starts, it must be processed continuously until completion, and its assigned machine cannot be changed during execution.
- Operations may span multiple consecutive shifts, which requires the assignment of multiple operators across the spanned shifts. Note that this assumption also includes operations with duration less than a shift, since they also may span two shifts.

- Operators are assigned to machines and shifts to ensure coverage of the scheduled operations.
- Operators are allowed to work at most one shift within any three consecutive shifts, as the shift timetabling regulation. Similar regulations have been deployed in related studies in the existing literature [Artigues et al., 2009].
- Only one operator is needed to operate a machine while it is processing an operation.
- The objective is to minimize the makespan of the schedule.
- The number of operators is limited, making the optimal makespan of the FJSP subproblem without operator-related restrictions infeasible.

To address this integrated complex problem, we propose a strengthened mixed-integer programming (MIP) formulation that significantly outperforms existing formulations in the literature [Ghorbani Saber, Leyman, and Aghezzaf, 2023]. On a set of generated test instances, the strengthened formulation improves the Linear Programming (LP) relaxation by an average of 13.16%, demonstrating its tighter representation of the feasible region.

Furthermore, we develop a Logic-Based Benders Decomposition (LBBD) framework in which the feasibility subproblem is formulated as an integer linear program. This approach generates valid feasibility cuts and guarantees optimality once no further cuts can be added. To enhance the computational efficiency and accelerate convergence, we incorporate cut strengthening techniques derived from an analysis of minimal infeasible sets. Moreover, based on a similar idea a matheuristic approach has been developed to address the problem. The computational results confirm that the strengthened LBBD and matheuristic methods outperform the standalone MIP formulation, offering improved scalability and solution quality for larger instances.

## References

- Christian Artigues, Michel Gendreau, Louis-Martin Rousseau, and Adrien Vergnaud. Solving an integrated employee timetabling and job-shop scheduling problem via hybrid branch-and-bound. *Computers & Operations Research*, 36(8):2330–2340, 2009. doi: 10.1016/j.cor.2008.08.013.
- Reza Ghorbani Saber, Pieter Leyman, and El-Houssaine Aghezzaf. Mixed-integer programming for integrated flexible job-shop and operator scheduling in flexible manufacturing systems. In *IFIP International Conference on Advances in Production Management Systems*, pages 460–470. Springer, 2023. doi: 10.1007/978-3-031-43688-8\_32.
- Olivier Guyon, Pierre Lemaire, Eric Pinson, and David Rivreau. Solving an integrated job-shop problem with human resource constraints. *Annals of Operations Research*, 213(1):147–171, 2014. doi: 10.1007/s10479-012-1132-3.

# Efficient Algorithms for Energy-Aware Single-Machine Scheduling with Battery Storage

L. Forte, S. De Vuyst, P. Leyman

Ghent University, Department of Industrial Systems Engineering and Product Design

Industrial Systems Engineering (ISyE), Flanders Make, Kortrijk, Belgium

luca.forte@ugent.be; stijjn.devuyst@ugent.be; pieter.leyman@ugent.be

## 1 Introduction

Energy-aware scheduling has become a central challenge in modern manufacturing environments. As industries increasingly aim to reduce operational costs and carbon emissions, efficiently coordinating production activities with electricity prices is more important than ever. Moreover, reducing electricity costs can have the desirable side effect of improving carbon emissions as well [2]. Introducing energy storage systems into the scheduling environment can further enhance cost savings; however, doing so significantly complicates the already usually complex problem structure. Despite its practical relevance, the algorithmic literature on energy-aware scheduling with storage systems remains limited: to the best of our knowledge, most existing works rely on mathematical modeling approaches, which quickly become computationally intractable for realistically sized instances.

## 2 Problem description

In this work, we study the problem of minimizing the energy cost of executing a set of jobs on a single machine within a fixed time horizon, where electricity prices follow a Time-of-Use (TOU) tariff. In addition, we consider the presence of an industrial battery that can be charged from the grid and discharged during high-price periods to reduce overall energy costs. Even without energy storage, this scheduling problem is NP-hard ([3], [1]); the introduction of a battery adds additional combinatorial interactions that further complicate the decision space.

## 3 Methodology

To tackle this problem, we begin by developing two mixed-integer linear programming models. The first, which we refer to as MILP-Schedule, computes the optimal job start times and battery usage for a fixed job sequence; although its scalability is limited, it provides an exact baseline against which all other approaches can be evaluated. The second model, MILP-Battery, determines the optimal battery usage given the complete job schedule and will be used as a

building block for metaheuristic approaches (see below).

To move beyond the scalability limitations of MILP-based approaches, we explore algorithmic strategies based on dynamic programming. We first design DP-Schedule, an exact pseudo-polynomial algorithm that produces the same optimal solutions as MILP-Schedule. Although exact, preliminary tests indicate that its running time is still too high to allow an exhaustive search over all job permutations, even for relatively small instances. We therefore introduce DP-Timing, a true polynomial-time algorithm that efficiently computes optimal job start times for any fixed ordering while ignoring the battery. This algorithm, together with MILP-Battery, becomes a key building block in more advanced strategies described below.

Building on these approaches, we outline several directions for developing approximate or metaheuristic algorithms capable of solving large instances. A first avenue is to employ metaheuristics that explore the space of job permutations while relying on DP-Timing and MILP-Battery to evaluate each candidate solution; additionally, variants that jointly search over job permutations and start time assignments will be investigated. Finally, another promising direction is to accelerate DP-Schedule by discretizing the range of possible battery levels, thereby trading optimality for speed and yielding an approximate but significantly faster algorithm.

Altogether, these ideas constitute the foundation for developing and comparing algorithms capable of producing high-quality solutions for realistically sized instances for energy-aware single machine scheduling with the inclusion of an industrial battery.

**Acknowledgement** We gratefully acknowledge the support provided by Ghent University’s Special Research Fund (BOF) under grant IDs BOF.BAF.2024.0972.01 and BOF.BAF.2025.0021.01.

## References

- [1] Kan Fang, Nelson A Uhan, Fu Zhao, and John W Sutherland. Scheduling on a single machine under time-of-use electricity tariffs. *Annals of Operations Research*, 238(1):199–227, 2016.
- [2] Nadine Gabrek and Stefan Seifermann. How the correlation between electricity prices and emission intensity affects the economic and ecological potential of industrial demand-side flexibility measures. *Journal of Cleaner Production*, page 145863, 2025.
- [3] Shijin Wang, Zhanguo Zhu, Kan Fang, Feng Chu, and Chengbin Chu. Scheduling on a two-machine permutation flow shop under time-of-use electricity tariffs. *International Journal of Production Research*, 56(9):3173–3187, 2018.

# Dynamic Programming for Inspection Station Allocation in Production Lines

Sajjad Hedayati<sup>1,2</sup>, Stijn De Vuyst<sup>1,2</sup>, Birger Raa<sup>1</sup>

<sup>1</sup> UGent, Dpt. Industrial Systems Engineering and Product Design

<sup>2</sup> Industrial Systems Engineering (ISyE), Flanders Make, Belgium

sajjad.hedayati@ugent.be, stijn.devuyst@ugent.be, birger.raa@ugent.be

## Introduction

In multi-stage manufacturing systems, the allocation of inspection stations is critical for minimizing total quality-related costs. This paper addresses the *Uncapacitated Inspection Station Allocation Problem with Delayed Inspection* (UISAPDI). This problem is a specialized variant of the *Inspection Station Allocation Problem with Delayed Inspection* (ISAPDI) introduced by Hedayati et al. [2]. Unlike traditional models [1] that enforce immediate inspection, the ISAPDI framework allows for *delayed inspection*. In our uncapacitated variant (UISAPDI), an inspection request  $i$  originating at a *release stage*  $r_i$  can be deferred up to the *final stage*  $n$ , creating an inspection interval  $[r_i, n]$ .

We formulate the problem as an Integer Linear Program (IP) and propose an exact backward-recursion Dynamic Programming (DP) algorithm. Computational results show the DP approach significantly outperforms standard IP solvers.

## Problem Formulation

Consider a serial production line with  $n$  stages ( $N = \{1, \dots, n\}$ ) and  $m$  inspection requests ( $R = \{1, \dots, m\}$ ). Each request  $i$  has a quantity  $q_i$ , a release stage  $r_i$ , and a unit penalty cost  $\pi_i$  if uninspected. Opening a station at stage  $k$  incurs setup cost  $s_k$ , and inspecting a unit of any request costs  $c_k$ . We define the binary variable  $Y_k \in \{0, 1\}$  to equal 1 if a station is allocated at stage  $k$ . The integer variables  $X_{ik}$  and  $W_i$  represent the number of units of request  $i$  that are inspected at stage  $k$  and are left uninspected, respectively. The objective (3) is to minimize the total cost. Constraints (2) ensure every unit is inspected or penalized. Constraints (3) link inspections to open stations.

$$\min \sum_{k \in N} s_k Y_k + \sum_{i \in R} \sum_{k=r_i}^n c_k X_{ik} + \sum_{i \in R} \pi_i W_i \quad (1)$$

$$\text{s.t. } \sum_{k=r_i}^n X_{ik} + W_i = q_i \quad \forall i \in R \quad (2)$$

$$X_{ik} \leq q_i Y_k \quad \forall i \in R, k \in \{r_i, \dots, n\} \quad (3)$$

## Dynamic Programming Approach

A fundamental property of the uncapacitated problem is the *non-splitting property*: due to linear costs, it is optimal to assign all units of a request  $q_i$  to a single station (or penalize all) rather than splitting them. This allows our DP to make decisions at the request level. We develop a backward DP where stage  $k$  represents the production stage. The value function  $F(k)$  is the minimum cost for requests released at stages  $r \geq k$ . The recursion decides to either *skip* stage  $k$  (incurring penalties) or *allocate* a station at a downstream stage  $\ell \in [k, n]$  to handle the request:

$$F(k) = \min \left\{ \begin{array}{l} \sum_{i:r_i=k} q_i \pi_i + F(k+1), \\ \min_{\ell=k, \dots, n} \left\{ s_\ell + \sum_{j=k}^{\ell} \sum_{i:r_i=j} q_i \min(\pi_i, c_\ell) + F(\ell+1) \right\} \end{array} \right\}$$

The base case is  $F(n+1) = 0$ . The complexity is  $O(m \cdot n^2)$ .

## Computational Results

We compared the DP against Gurobi [5] (IP) on instances up to  $N = 100$ . Both methods found identical optima ( $Z_{IP} = Z_{DP}$ ). However, the DP was significantly faster, as detailed in Table 1. For  $N = 100$ , the DP was  $\approx 8.6\times$  faster than the IP.

Table 1: Average runtimes (seconds) and speedup for varying stages  $N$ .

$N$	IP (s)	DP (s)	Speedup	$N$	IP (s)	DP (s)	Speedup
10	0.0028	0.0005	11.45	60	0.1099	0.0128	7.82
20	0.0089	0.0012	8.54	70	0.1582	0.0178	8.57
30	0.0212	0.0029	8.02	80	0.2170	0.0248	8.73
40	0.0396	0.0055	7.23	90	0.2850	0.0320	8.77
50	0.0730	0.0087	7.47	100	0.3655	0.0394	8.60

## References

- [1] Lindsay, G.F., 1964, Optimal inspection policy for a serial production process. *PhD Thesis, University of California, Berkeley*.
- [2] Hedayati, S., De Vuyst, S., Raa, B., 2022, Allocation of inspection stations in a multi-product serial assembly line with delayed inspection. *FEARS 2022 Symposium*.

- [3] Bellman, R., 1966, Dynamic Programming. *Science*, 153(3731), 34-37.
- [4] Shiau, Y.R., 2002, Inspection allocation planning for a multiple-stage manufacturing system with inspection error. *Computers & Industrial Engineering*, 42(2), 205-216.
- [5] Gurobi Optimization, LLC, 2025, Gurobi Optimizer Reference Manual. <http://www.gurobi.com>.

# An integrated approach to operational scheduling in hospital-at-home

Anisha Maharani, Véronique François, Yasemin Arda

University of Liège, HEC Liège Management School

QuantOM, Research Centre for Quantitative methods and Operations Management

e-mails: {anisha.maharani, veronique.francois, yasemin.arda}@uliege.be

Recent advances in digital health technologies have driven the emergence of hospital-at-home (HaH) programmes over the past decades. HaH programmes deliver short-term, hospital-level treatments in patients' homes for conditions that would otherwise require inpatient hospitalization. The benefits of such model ranging from improved hospital resource utilization to enhanced patient comfort and well-being [4].

HaH admissions depend not only on the medical conditions and home environment of each patient, but also on the hospital capacity. One of the factors constraining resource capacity is the work schedules, referred to as rosters, of hospital nurses responsible for delivering HaH services. These rosters are typically constructed several months in advance, taking into account anticipated patient demand and nurse availability, as well as legal constraints and contractual agreements. As real-time information on patient requests and nurse availability becomes available, the initial roster can be adjusted to allocate HaH resources more effectively and maintain a high level of service. In this work, the process of updating the baseline roster, referred to as rerostering, is triggered by the arrival of new patients at the beginning of each week.

Our study addresses the operational level complexities of HaH scheduling by simultaneously making the following weekly decisions: selecting eligible patients for HaH admission, adjusting the roster as needed, assigning nurses to specific care tasks, and scheduling daily care visits for each nurse. Specifically, the present work integrates the home health care routing and scheduling problem [1] and the nurse rerostering problem [2], giving rise to the nurse routing and rerostering problem. The objective of the problem is defined in lexicographic order: it first maximizes the number of admitted patients to HaH, and then minimize the total routing duration of all nurses. The routing subproblem explicitly considers the compatibility between nurses' skills and the skills required by patients. When rerostering is necessary, adherence to legal and contractual constraints must be preserved. In our work, deviations from each nurse's schedule in the baseline roster are limited. Furthermore, this study implements a stepping horizon approach [3] to maintain continuity in respecting rostering constraints and in delivering the treatment for ongoing patients.

Two neighborhood search-based solution algorithms are specifically designed to address the interplay among different decisions in this rich and highly integrated problem. During the course of the algorithms, infeasible solutions are

allowed to be visited by relaxing time window, rostering, and rostering constraints. Both algorithms have the same structure consisting of an outer loop that modifies the set of admitted patients and an inner loop that improves the performance of rostering and routing subproblems. The main difference between the proposed algorithms lies in how patient set is constructed and modified. Destroy-and-repair mechanisms, inspired by large neighborhood search, are employed for modifying the patient set and refine routing-related performance. Classical local search operators are employed to restore feasibility with regard to rostering and rostering constraints.

The talk will first introduce new benchmark instances specifically built to represent typical HaH cases. It will then present numerical results that demonstrate the effectiveness and performance of the proposed solution approaches. Finally, insights into the impact of several instance parameters will be discussed, with a particular focus on treatment characteristics that enable the admission of more patients to the HaH system.

## References

- [1] Castillo-Salazar, J. A., Landa-Silva, D., & Qu, R. (2016). Workforce scheduling and routing problems: literature survey and computational study. *Annals of Operations Research*, 239(1): 39-67.
- [2] Moz, M., & Pato, M. V. (2003). An integer multicommodity flow model applied to the rostering of nurse schedules. *Annals of Operations Research*, 119, 285–301. Springer.
- [3] Salassa, F., & Vanden Berghe, G. (2012). A stepping horizon view on nurse rostering. In *Proceedings of the 9th international conference on the practice and theory of automated timetabling*. (pp. 161-173).
- [4] Shepperd, S., & Iliffe, S. (2005). Hospital at home versus in-patient hospital care. *Cochrane database of systematic reviews*, (3).

# Periodic home healthcare routing with stochastic patients : imitation learning enriched decision-making

Louise Tassin

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative methods and Operations Management  
e-mail: [louise.tassin@uliege.be](mailto:louise.tassin@uliege.be)

Véronique François

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative methods and Operations Management  
e-mail: [veronique.francois@uliege.be](mailto:veronique.francois@uliege.be)

Michel Gendreau

Polytechnique Montréal, Department of Mathematics and Industrial Engineering  
CIRRELT, Interuniversity Research Centre on  
Enterprise Networks, Logistics and Transportation  
e-mail : [michel.gendreau@polymtl.ca](mailto:michel.gendreau@polymtl.ca)

Elise Vandomme

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative methods and Operations Management  
e-mail: [elise.vandomme@uliege.be](mailto:elise.vandomme@uliege.be)

Yasemin Arda

University of Liège, HEC Liège Management School  
QuantOM, Research Centre for Quantitative methods and Operations Management  
e-mail: [yasemin.arda@uliege.be](mailto:yasemin.arda@uliege.be)

## 1 Introduction and problem definition

Home healthcare (HHC) is gaining popularity as a means of alleviating pressure on healthcare systems and improving patient well-being. However, effective HHC operations require optimized resource utilization, a challenge complicated by inherent uncertainty and dynamism. Although deterministic variants of the home healthcare routing and scheduling problem have been extensively studied, this is not yet the case of stochastic and dynamic versions, which are particularly challenging [Khorasanian et al., 2024].

Aiming to address this gap, our work analyzes the problem of a home healthcare provider (HHCP) making daily decisions regarding the acceptance of new

patient requests that arrive dynamically over the planning horizon. The HHCP also has to assign the first visit of each accepted patient to a day. It is assumed that every patient needs at least one visit and a specific periodicity of care, but the total number of visits that a patient will effectively require is not known with certainty. The HHCP must commit to serving accepted patients during their whole episode of care, while creating routes for the available caregivers in order to cover the visits planned for the current day. In case of excessive workload due to poor planning decisions, the HHCP can call on external nurses as daily recourse actions.

The objective of the HHCP is to accept as many patient requests as possible and to minimize its total operating cost, including routing and outsourcing costs, over the planning horizon.

## 2 Solution method

We model this problem as a Markov decision process and propose a solution approach that integrates machine learning techniques with traditional combinatorial optimization algorithms. This hybrid approach aims to tackle the stochasticity inherent in the patients arrival process and their care durations using a neural network (NN).

The NN will be trained on "perfect information" solutions in an imitation learning fashion [Dalle et al., 2022, Baty et al., 2024]. This training dataset is composed of solutions to the static and deterministic version of the problem of interest, obtained from historical data. This problem variant is a periodic home healthcare routing problem, and can therefore be solved using an adaptation of the tabu search algorithm proposed by Cordeau et al. [Cordeau et al., 1997] for the periodic vehicle routing problem. The expertise gained by the neural network is then embedded into the same tabu search algorithm, which will be employed on a rolling-horizon basis in order to develop a decision-making process for stochastic and dynamic instances.

## References

- [Baty et al., 2024] Baty, L., Jungel, K., Klein, P. S., Parmentier, A., and Schiffer, M. (2024). Combinatorial optimization-enriched machine learning to solve the dynamic vehicle routing problem with time windows. *Transportation Science*, 58(4):708–725.
- [Cordeau et al., 1997] Cordeau, J.-F., Gendreau, M., and Laporte, G. (1997). A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30(2):105–119.
- [Dalle et al., 2022] Dalle, G., Baty, L., Bouvier, L., and Parmentier, A. (2022). Learning with combinatorial optimization layers: a probabilistic approach.
- [Khorasanian et al., 2024] Khorasanian, D., Patrick, J., and Sauré, A. (2024). Dynamic home care routing and scheduling with uncertain number of visits per referral. *Transportation Science*, 58(4):841–859.

## Merging HIS- and RTLS-data to Strengthen Work Organization Analysis in Nursing

Haroon Tharwat

UHasselt - Hasselt University, Digital Future Lab

e-mail: [haroon.tharwat@uhasselt.be](mailto:haroon.tharwat@uhasselt.be)

Niels Martin

UHasselt - Hasselt University, Digital Future Lab

e-mail: [niels.martin@uhasselt.be](mailto:niels.martin@uhasselt.be)

Benoit Depaire

UHasselt - Hasselt University, Digital Future Lab

e-mail: [benoit.depaire@uhasselt.be](mailto:benoit.depaire@uhasselt.be)

Nursing work is under pressure due to rising care demands and limited staffing. Hiring more nurses is costly and challenging, which makes it essential to improve their work organization. To identify improvement areas, analyzing how nurses currently work can be very insightful. This study will use process data for work organization analysis that does not put an extra registration burden on nurses.

In most hospitals, the primary source for recording nursing tasks is the Hospital Information System (HIS). However, records in HIS-data can deviate from actual nursing work: several tasks are not recorded, and those that are recorded may include incomplete or inaccurate information. For example, documentation times can differ significantly from the actual execution time. Such discrepancies create gaps and delays that render work organization analysis less reliable when HIS data is used in isolation (Vanthienen et al. [2025]). To address these limitations, alternative sources, such as Real-Time Location System (RTLS) data, have been explored (Fernández-Llatas et al. [2015]). RTLS traces reveal when resources enter or exit specific locations. However, they do not indicate which task was performed. As both HIS- and RTLS-data provide a partial view of nursing work, investigating their integrated use is worthwhile to obtain a more accurate view of task start and end times, locations, and resources (Martin [2018]).

Because integrating HIS- and RTLS-data is non-trivial, this study proposes FINTR, a Footprint-based Integration method for Nursing Task Retrieval. FINTR is a semi-automated approach that merges the two data sources using task footprints. A task footprint encodes domain knowledge and expresses how the execution of a type of nursing task is reflected in HIS- and/or RTLS-data. FINTR uses these footprints to guide and constrain the data integration process.

FINTR consists of three main steps. It first performs candidate filtration to remove combinations of HIS entries and RTLS intervals that cannot represent the same task. For the remaining possibilities, it then scores each candidate based on how well it fits the expected temporal and structural patterns encoded in the

task footprint. Finally, FINTR selects the most plausible candidate and merges its HIS and RTLS evidence into a single constructed task instance. Thus, FINTR produces a nursing task log where each entry represents a performed nursing task.

Preliminary evaluations were performed using synthetic data, which has the advantage that the ground truth nursing task log is known, to which the output of FINTR can be compared. Two simulation settings have been considered so far: one in which the documentation times in HIS correspond to the moment of task execution and another in which there are documentation delays of up to ten minutes (i.e., tasks are documented in the HIS after they have been performed). Initial results suggest that FINTR performs well, but further testing is needed.

Overall, the findings show that FINTR offers a viable approach to integrating HIS- and RTLS-data to construct a nursing task log from routinely collected digital traces. The method provides a systematic, reusable framework for creating more reliable logs to support downstream analyses. By offering a more accurate and detailed representation of nursing work, FINTR opens new possibilities for understanding and improving the organization of care.

## References

- C. Fernández-Llatas, A. Lizondo, E. Montón, J. M. Benedi, and V. Traver. Process mining methodology for health process tracking using real-time indoor location systems. *Sensors*, 15(12):29821–29840, 2015.
- N. Martin. Using indoor location system data to enhance the quality of healthcare event logs: Opportunities and challenges. In *International Conference on Business Process Management*, pages 226–238. Springer International Publishing, 2018.
- A. Vanthienen, N. Martin, and B. Depaire. Investigating the (mis)match between electronic health records and actual nursing work: An observational study. *International Journal of Nursing Studies*, 151:105309, 2025.

## Integrated planning of public charging infrastructure and low-voltage grid expansion under uncertainty

Bryan Coulier

KU Leuven, Department of Computer Science

e-mail: [bryan.coulier@kuleuven.be](mailto:bryan.coulier@kuleuven.be)

Hatice Çalık

KU Leuven, Department of Electrical Engineering

Thijs Becker

Flemish Institute for Technological Research (VITO),

Greet Vanden Berghe

KU Leuven, Department of Computer Science

Rapidly evolving targets and forecasts for electric vehicle adoption create deep uncertainty for the planning of urban public charging infrastructure. Effective planning must ensure that the low-voltage grid, which integrates these charging stations, remains capable of safely handling increased electricity demand. Currently, decisions regarding public charging station locations and low-voltage grid expansions are often made by separate stakeholders whose interests overlap but nevertheless do not interact. This lack of coordination can lead to grid-blind placement of chargers or costly, poorly targeted grid reinforcements. This study addresses these challenges by proposing an integrated planning methodology for public charging station location and low-voltage grid expansion. The methodology assists public sector planners and distribution system operators in making informed, resilient decisions despite inherent future uncertainties. The proposed methodology explicitly models uncertainty in public charging demand and baseline household loads, avoiding reliance on a single forecast. We quantify the trade-off between investment costs and grid impact and identify substations that remain persistently problematic across plausible future scenarios. By prioritizing investments in substations with persistent bottlenecks the approach enables proactive infrastructure deployment.

# The Dial-a-Ride Problem with Synchronized Visits

Boshuai Zhao

ORSTAT, Faculty of Economics and Business, KU Leuven  
and Université Paris-Saclay, CentraleSupélec, LGI  
e-mail: [boshuai.zhao@kuleuven.be](mailto:boshuai.zhao@kuleuven.be)

Jakob Puchinger

EM Normandie Business School  
and Université Paris-Saclay, CentraleSupélec, LGI  
e-mail: [jakob.puchinger@centralesupelec.fr](mailto:jakob.puchinger@centralesupelec.fr)

Roel Leus

ORSTAT, Faculty of Economics and Business, KU Leuven  
e-mail: [roel.leus@kuleuven.be](mailto:roel.leus@kuleuven.be)

The limited capacity of drones and future one- or two-seat modular vehicles requires multiple units to serve a single large customer (i.e., a customer whose demand exceeds a single vehicle's capacity) simultaneously, whereas small customers (i.e., those whose demand can be served by a single vehicle) can be consolidated in one trip. This motivates the Dial-a-Ride Problem with Synchronized Visits, where a fleet of drones must be routed and scheduled to transport orders at minimum cost. We propose four formulations: arc-based, event-based, time-space event-based (TSEF), and time-space fragment-based (TSFrag). An event is defined as a tuple of a location and a set of onboard customers, while a fragment represents a partial path. For TSEF and TSFrag, we also employ the dynamic discretization discovery (DDD) algorithm [1], which iteratively refines an initial low-resolution time-space network to obtain a continuous-time optimal solution. Computational results show that the event-based formulation performs best under low request intensity (few customers per unit time), whereas TSFrag with DDD excels with high request intensity; both substantially outperform the arc-based formulation. When implemented with DDD, TSFrag also requires less time and fewer iterations than TSEF. We also apply our methods to the classical dial-a-ride problem, where we find that that TSFrag with DDD can replace call-backs in case of high request intensity, and that using DDD is more beneficial to this problem than to the pickup-and-delivery problem with time windows.

## References

- [1] Boland N, Hewitt M, Marshall L, Savelsbergh M (2017) The continuous-time service network design problem. *Oper. Res.* 65(5):1303–1321.

# Fast Passenger Assignment for Public Transport Services

Maarten Wens

KU Leuven, Institute for Mobility - CIB

e-mail: [maarten.wens@kuleuven.be](mailto:maarten.wens@kuleuven.be)

Jeroen Verstraete

KU Leuven, Institute for Mobility - CIB

Pieter Vansteenwegen

KU Leuven, Institute for Mobility - CIB

In public transport planning, properly modelling how passengers will use a proposed transport service is crucial. This aspect of the planning process, known as the Passenger Assignment Problem (PAP), involves modelling passenger behaviour on a transport network and poses significant computational challenges [1]. The PAP must be solved repeatedly during many planning processes [2]. A common approach is to calculate all-to-all shortest paths on the network, which are then combined with passenger demand to estimate average travel times—a key metric in transportation planning.

To address this challenge, we propose a novel transit network representation that enables the decomposition of the original networks into smaller networks, thereby reducing redundant calculations and the required computational time for solving the PAP. This structure enables existing algorithms to run faster and more efficiently. The core idea is to decompose the network into several subgraphs that contain only the nodes that recur in many partial paths. An example of such a partial network is a subgraph containing only transfer nodes, where passengers can change lines, or a node placed on a bridge separating two parts of the network. The complete path of the traveller can then be split up into a path from an origin to a subgraph, a path over the subgraph and a path from the subgraph to the destination. As a result, the paths in (and through) the subgraph only need to be calculated once. This method is furthermore recursive and allows for breaking down the initial network into several layers.

The main problem with solving the PAP in a transit network is that oftentimes a cost is incurred at the nodes of the graph, either through dwell or transfer times. To overcome this difficulty, a graph transformation is necessary, after which a shortest path algorithm can be executed. In the current state of the art, initial graph transformations, such as the change&go network [3], add dummy nodes between which the transfer times can be modelled as an arc cost. This increases the number of nodes in the network significantly, and hence the calculation time.

Instead of adding additional nodes, Aktaş, Vermeir, and Vansteenwegen [1]

introduce a network transformation that introduces only additional edges, no additional nodes, called the Direct Link Network (DLN). We expand on this concept by showing how this network representation can be adapted to include dwell times, frequencies, and even complete time tables. In the latter case, we consider periodic time tables, which allow for a more compact graph version of the network compared to alternative representations, and show an efficient way to perform an assignment of all passengers on this network.

When this network transformation is combined with the network decomposition, the PAP for many real-sized instances can be calculated in a negligible amount of time. e.g. A rural network containing 500 stops and 70 distinct buslines can be solved in  $\sim 0.01$ s when every transfer is penalised with a general transfer penalty. With the DLN approach, it would take  $\sim 0.1$  seconds to calculate the same result, an additional order of magnitude. With our decomposition, the same network can be solved in  $\sim 0.34$ s when the problem is extended and waiting times are modelled as a function of the frequency of the departing line, and lines are given dwell times. When detailed timetables are included, the same instance takes  $\sim 2.74$ s to perform a full passenger assignment. In this last case, the transfer time depends on the time difference between arriving and departing vehicles, while the initial waiting time for the first vehicle depends on the difference between the desired moment of departure and the scheduled departures. When timetables are included, individual vehicles with unique driving and waiting times need to be modelled, along with individual passenger requests. Because of the additional complexity, the increase in calculation time seems reasonable.

Future research can determine whether a more efficient implementation of a shortest path algorithm, or a combination with other state-of-the-art techniques, can lead to further reduction in calculation times. We consider this research a significant step forward in showing the increase in complexity when passenger assignment is performed with increasing levels of detail.

## References

- [1] Aktaş, Dilay, Evert Vermeir, and Pieter Vansteenwegen. "Speeding up the Passenger Assignment Problem in transit planning by the Direct Link Network representation." *Computers & Operations Research* 167 (2024): 106647.
- [2] Iliopoulou, Christina, Konstantinos Kepaptsoglou, and Eleni Vlahogianni. "Metaheuristics for the transit route network design problem: a review and comparative analysis." *Public Transport* 11 (2019): 487-521.
- [3] Schöbel, Anita. "Line planning in public transportation: models and methods." *OR spectrum* 34.3 (2012): 491-510.

## Column elimination for scheduling problems

Vianney Coppé

Hexaly

e-mail: [vcoppe@hexaly.com](mailto:vcoppe@hexaly.com)

Column elimination [2] is a recent framework for solving combinatorial optimization problems in which the solution consists of multiple sequences. It uses a *relaxed decision diagram* to compactly encode a superset of all feasible sequences. By solving a constrained network flow problem in this decision diagram (DD), a (relaxed) minimum-cost set of sequences can be found. If it happens to correspond to an exact solution, it is also guaranteed to be the optimal one. Otherwise, this relaxed solution gives a lower bound on the optimal cost, and the DD can be refined to *eliminate* the relaxed solution at hand before solving a new iteration of the flow problem. This technique was successfully applied to several vehicle routing and graph coloring problems [2, 3]. In this paper, we investigate how to apply it to disjunctive scheduling problems.

There are two main obstacles to applying column elimination to scheduling problems. First, while routing problems typically involve uniform vehicles or few vehicle classes, the machines considered in scheduling problems are usually all different. This is due to heterogeneous machine-task compatibilities and possibly machine-dependent task durations, setup and changeover times. We can deal with this variety of machines by creating a separate relaxed DD for each of them and embedding those into a single *meta-DD*. This meta-DD simply connects its root node to each machine-specific root node, and each machine-specific terminal node to its own terminal node. Since machine-task compatibilities are generally sparse, the size of the meta-DD remains reasonable.

The second obstacle lies in the presence of precedence constraints between tasks. In routing problems, it is always preferable to visit clients as early as possible. However, scheduling problems with precedence constraints may require machines to wait for a task to be completed on another machine before starting a given task. We first explain how to allow such idle times and then how to enforce precedence constraints. To model idle times, the DD must include transitions for a whole range of start times for each task. Clearly, representing all possible time-indexed transitions would yield very large DDs, especially when the horizon is loose. We propose to use the modeling introduced in [1] for single-machine constraint propagation in constraint programming, in which time information is not explicitly encoded in the DD but rather derived a posteriori. It works by performing a top-down and bottom-up pass on the DD to compute information such as the earliest and latest completion times for each node  $n$ , respectively denoted  $ect(n)$  and  $lct(n)$ . This information is used to identify and filter infeasible transitions, and to assign transition costs to the remaining arcs.

Because they are computed with respect to the most favorable completion time, the transition costs may be relaxed. Therefore, after finding a minimum-

cost flow in the DD, we compute the true completion time of each task in the solution and use it to obtain the true transition costs. If all the transition costs match, the solution is exact. Otherwise, we identify transitions having a relaxed cost and refine them as follows. Let  $a = (u, v)$  be an arc with relaxed cost and let  $ct_a$  be the true completion time of the corresponding transition in the solution. We refine the DD by splitting node  $v$  into two nodes  $v_1$  and  $v_2$  with shorter completion time intervals. More precisely, if  $ect(v) < ct_a \leq lct(v)$ , we create nodes  $v_1$  and  $v_2$  with explicitly encoded completion time intervals  $[ect(v_1), lct(v_1)] = [ect(v), ct_a - 1]$  and  $[ect(v_2), lct(v_2)] = [ct_a, lct(v)]$ . We call this operation *splitting* because all the incoming and outgoing arcs of  $v$  are then transferred to nodes  $v_1$  and  $v_2$  before removing node  $v$ . After a certain number of such refinements, a new round of propagation and filtering can be performed to filter additional infeasible transitions and improve the arc costs.

Let us now explain how to enforce precedence constraints. We propose to handle them by embedding column elimination in a *branch-and-bound* (B&B) framework. After applying column elimination in a given B&B node, we look for violated precedence constraints in the solution. If there are none, the solution is feasible, and may later be proven optimal by the B&B. Otherwise, let  $i$  and  $j$  be two tasks with start and completion times  $st_i, ct_i$  and  $st_j, ct_j$  in the solution, such that  $i$  must precede  $j$  but  $ct_i > st_j$ . We create a first branch with the additional constraint  $st(j) < ct_i$  and a second with  $st(j) \geq ct_i$ . Those time constraints can be propagated through the precedence network to tighten the earliest and latest times for each task. They are then injected into column elimination by finding all transitions that are incompatible with them and blocking any flow through those. Note that branching information can also be used to strengthen the arc costs and to determine the true completion times inside the refinement procedure.

It may happen that a combination of branching constraints creates an infeasible subproblem. To detect those cases, we apply column elimination on the linear relaxation of the constrained network flow problem and solve it with a linear programming solver. Therefore, we may also need to branch to separate fractional solutions – a procedure called *branch-and-refine* in [2] – using task time constraints and forced/forbidden machine-task assignment constraints.

This approach was implemented inside the Hexaly global optimization solver and significantly improves the bounds computed for several variants of the (flexible) job shop scheduling problem, increasing the proportion optimality proofs on the benchmark instances. It can handle all standard scheduling objectives; those including a *maximum* function need to use a slightly adapted flow problem.

## References

- [1] Andre A. Cire and Willem-Jan van Hoeve. Multivalued decision diagrams for sequencing problems. *Operations Research*, 61(6):1411–1428, 2013.
- [2] Anthony Karahalios and Willem-Jan van Hoeve. Column elimination for capacitated vehicle routing problems. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 35–51. Springer, 2023.
- [3] Willem-Jan van Hoeve. Graph coloring with decision diagrams. *Mathematical Programming*, 192(1):631–674, 2022.

# Integrating Public Transportation in Bike-Sharing Network Design: A Mathematical Model and Numerical Study

Soukaina Bayri

Hasselt University, Research Group Logistics  
Vrije Universiteit Brussel, Research Group Mobilise  
e-mail: [soukaina.bayri@uhasselt.be](mailto:soukaina.bayri@uhasselt.be)

Kris Braekers

Hasselt University, Research Group Logistics  
e-mail: [kris.braekers@uhasselt.be](mailto:kris.braekers@uhasselt.be)

An Caris

Hasselt University, Research Group Logistics  
e-mail: [an.caris@uhasselt.be](mailto:an.caris@uhasselt.be)

Over the past few years, Bike-Sharing Systems (BSSs) have been widely adopted and have become a familiar part of the urban landscape. These systems offer numerous benefits for both the user and the environment, such as reduced emissions, improved public health, and less traffic congestion. In practice, the stations where shared bikes are parked are often strategically placed near Public Transportation (PT) stops to expand the coverage of the PT network and to facilitate seamless connections between BSSs and PT. The success of these multi-modal systems combining bike-sharing and PT, heavily depends on the locations of the bike-sharing stations in relation to the PT network [2]. Bike-sharing stations that are poorly positioned in relation to the PT network will compromise the success of the overall system [1][4]. However, the literature on bike-sharing planning problems in a multi-modal context is very limited, and lacks methodologies to support decisions on BSSs while taking into account other transportation modes [3][5][6]. Most papers study BSSs in an isolated context without considering other modes, which is the gap we address with our research.

In this talk, we present a novel optimization problem and a corresponding mathematical model for the design of BSSs that complement existing PT networks. The model determines the optimal locations and capacities of bike-sharing stations in relation to the PT network, and the corresponding bike repositioning requirements. This model is then used in a numerical study, to quantify the impact of incorporating the PT network during the planning phase. The impact is evaluated across a variety of scenarios, including differing PT network densities and both random and clustered demand distributions. Comparing results across different scenarios enables us to assess the added value of integrating PT information in the planning phase under diverse network structures. Our findings

provide methodological and managerial insights for the design of BSSs within a multi-modal mobility context.

## References

- [1] Frade, I., & Ribeiro, A. (2015). Bike-sharing stations: A maximal covering location approach. *Transportation Research Part A: Policy and Practice*, 82, 216-227.
- [2] Lin, J. R., & Yang, T. H. (2011). Strategic design of public bicycle sharing systems with service level constraints. *Transportation research part E: logistics and transportation review*, 47(2), 284-294.
- [3] Shui, C. S., & Szeto, W. Y. (2020). A review of bicycle-sharing service planning problems. *Transportation Research Part C: Emerging Technologies*, 117, 102648.
- [4] Bahadori, M. S., Gonçalves, A. B., & Moura, F. (2021). A systematic review of station location techniques for bicycle-sharing systems planning and operation. *ISPRS international journal of geo-information*, 10(8), 554.
- [5] Caggiani, L., Colovic, A., & Ottomanelli, M. (2020). An equality-based model for bike-sharing stations location in bicycle-public transport multimodal mobility. *Transportation Research Part A: Policy and Practice*, 140, 251-265.
- [6] Zhang, C., & Zhou, H. (2019). The study of coepetition between public bus and bike sharing based on environmental protection. In *E3S Web of Conferences* (Vol. 136, p. 04015). EDP Sciences.

## ORBEL Award session

**Candidates:**

**Jitse De Latte** Integration of human factors into an Assembly Line-feeding decision support model focusing on fatigue and cognitive load

This thesis studies the integration of human factors in Assembly Line Feeding (ALF) decision support. For the purposes of this investigation, an existing MIP model developed for Tactical ALF optimization is extended to consider operator fatigue, cognitive load and ergonomic constraints. A cross-evaluation experimental analysis is conducted to assess the performance of the original and the extended model in each other's context. Results show how the human factor-aware model may significantly reduce physical and cognitive burden, enhancing worker well-being and supporting sustainable production with only a marginal monetary cost increase.

**Quentin Reussens** Project Interdiction with Stochastic Activity Durations

This thesis addresses project interdiction under uncertainty, an optimization problem where an adversary strategically disrupts activities to maximally delay a project within budget constraints. By comparing three formulations (deterministic, exponential, and hypo-exponential models), the work reveals how different uncertainty representations change optimal interdiction strategies. The work introduces hypo-exponential distributions to better capture variability inside activities, identifies key features driving interdiction decisions (duration increase, budget cost, and slack time) and demonstrates the importance of accurate intelligence in competitive environments. To address computation complexity, four polynomial-time heuristics are developed that achieve near-optimal performance in small networks while reducing solution times, making the approach more practical than the computation of optimal solutions.

**Anne-Sophie Marx and Lynn Goossens** Optimizing Urban Pedestrian : Bi-Level Network Design Approach to Pedestrian Zone Planning and Vehicular Impact Analysis

This thesis examines the trade-offs between pedestrianization and vehicular travel time within urban street networks by formulating a bi-level Network Design Problem (NDP). To address the competing objectives of minimizing total travel time, maximizing pedestrian network connectivity, and maximizing the total length of pedestrianized streets, a mixed-integer nonlinear bi-level programming model is proposed where the upper level models the decisions of a central planner selecting streets to close to motorized traffic, while the lower level simulates the behavior of vehicle users in response to upper-level decisions. The model is applied to the Sioux Falls network. Results suggest that a carefully balanced pedestrianization strategy can lead to pedestrianized street configurations with minimal impact on total travel time and indicate that an optimal level of pedestrian network connectivity does not necessarily correspond to the optimal level of pedestrianized length, and vice versa.

# Integrating stochastic optimization and decision-focused learning

Alessandro Barbini

Ghent University, KERMIT, BIOVISM

e-mail: [alessandro.barbini@ugent.be](mailto:alessandro.barbini@ugent.be)

Bernard De Baets

Ghent University, KERMIT

e-mail: [bernard.debaets@ugent.be](mailto:bernard.debaets@ugent.be)

Jan Verwaeren

Ghent University, BIOVISM

e-mail: [jan.verwaeren@ugent.be](mailto:jan.verwaeren@ugent.be)

Consider a shortest-path problem where the edge weights are not directly available. Instead, the decision maker is given properties that are correlated with the edge weights and can be used as inputs to a machine-learning model to predict them. In a second step, the decision maker can use these predicted weights to compute the shortest path via a solver. This setup illustrates the classic *predict-then-optimize* (PTO) approach to this type of problems. Despite its intuitive nature, it has been shown [1] that predictive accuracy does not necessarily translate into high-quality decisions (in this case, finding the *true* shortest path). Decision-Focused Learning (DFL) addresses this limitation by training models end-to-end using loss functions that compare the quality of solutions obtained with predicted edge weights with the true edge weights [1]. As a result, the prediction models learn to predict edge weights that lead to high-quality solutions.

In this work, we extend the DFL framework to stochastic optimization problems, where the parameters of the optimization problem are (possibly dependent) random variables. In such cases, the objective often is to optimize the expected, median quality, or even worst case quality of a solution. In analogy with the traditional DFL setting, we consider situations where the underlying distributions are directly known but can be estimated using correlated properties. We refer to this setting as *Probabilistic Decision Focused Learning*. This framework closes the gap between DFL and stochastic optimisation by learning probability distributions that are explicitly optimised to improve decision quality. It is particularly advantageous when structural dependencies and uncertainty play a significant role in determining outcomes.

Our primary contribution is to show that structured probabilistic models, such as conditional random fields (CRFs) [2], can be trained end-to-end within a DFL objective, substantially improving decision quality. Specifically, we model dependencies using a linear-chain graphical model in which the nodes represent items and edges encode local dependencies. We compute gradients with respect to the optimal solutions using existing gradient approximation techniques such as I-MLE [3] and Perturb-and-MAP [4]. As a result, the loss function depends directly on the solutions themselves rather than on discrepancies between the estimated and true CRF parameters.

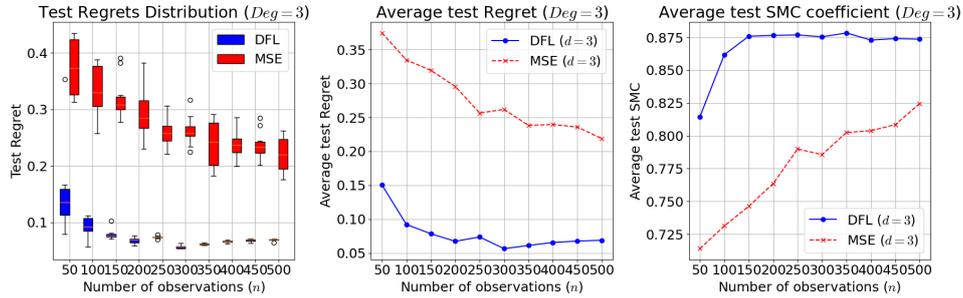


Figure 1: DFL and MSE performance as a function of training set size from (50-500) for a degree-5 chain data-generation process. The left panels show box-plots of relative regret; the middle panels show the mean relative regret across the 10 models; and the right panels show the mean SMC coefficient across the models.

We evaluate our approach using both synthetically generated datasets [5] and a 20-year rainfall time-series dataset [6]. In synthetic experiments with linear predictive models, probabilistic DFL consistently outperforms PTO, achieving roughly a 20% reduction in regret compared with the standard PTO pipeline (see Figure ). This performance gap becomes more pronounced as the data-generation process increases in complexity. In simpler scenarios, the standard approach eventually catches up with DFL when very large amounts of data are available. However, such convergence is no longer observed in more complex settings.

## References

- [1] Jayanta Mandi, Peter J Stuckey, Tias Guns, et al. Smart predict-and-optimize for hard combinatorial optimization problems. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 1603–1610, 2020.
- [2] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1529–1537, 2015.
- [3] Mathias Niepert, Pasquale Minervini, and Luca Franceschi. Implicit mle: backpropagating through discrete exponential family distributions. *Advances in Neural Information Processing Systems*, 34:14567–14579, 2021.
- [4] George Papandreou and Alan L Yuille. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *2011 international conference on computer vision*, pages 193–200. IEEE, 2011.
- [5] Adam N Elmachtoub and Paul Grigas. Smart “predict, then optimize”. *Management Science*, 68(1):9–26, 2022.
- [6] Meteostat. Meteostat: Weather and climate data. <https://meteostat.net>, 2025.

# Temporal Model-Agnostic Meta-Learning for Cold-Start Time Series Forecasting

Wannes Janssens, Matthias Bogaert, Dirk Van den Poel

Ghent University, Department of Marketing, Innovation and Organisation

FlandersMake@UGent-corelab CVAMO, Ghent, Belgium

e-mail: [wanjanss.janssens@ugent.be](mailto:wanjanss.janssens@ugent.be)

**Keywords:** Cold-Start, Meta-Learning, MAML, Time Series Forecasting

Time series forecasting is a foundational element of data-driven decision-making across domains such as supply chain management, retail, energy, and healthcare. In recent years, methodological advances, as well as the expanding availability of large-scale time series datasets, have increasingly favoured deep learning-based forecasting approaches [1]. The cold-start problem arises when a time series contains little or no historical data, making accurate forecasting particularly challenging. This scenario is especially relevant in manufacturing contexts, such as predicting demand for new parts, where historical sales data may be limited or nonexistent.

In the literature, two main approaches have been proposed to address the cold-start problem. The first estimates the similarity between training and test instances using meta-features (e.g. product attributes) to produce cold-start forecasts [2]. The second uses global forecasting models, which can incorporate meta-features and are trained on collections of related time series but are optimized for series seen during training. While such models can reasonably generalize to unseen series when trained on large datasets, their cold-start performance in small-data regimes remains limited. Moreover, both approaches rely on the availability of rich and informative attributes, which is often not guaranteed.

Model-Agnostic Meta-Learning (MAML) [3] provides a general framework for rapid adaptation to new tasks and has demonstrated strong few-shot performance in domains such as computer vision, recommender systems, and time series forecasting. In existing time series applications, tasks are typically defined at the entity (e.g. a consumer) or dataset level [4], with each task comprising multiple related series. In contrast, we define each task as a single demand series and study temporal cold-start forecasting, where adaptation must occur from only a short initial history rather than across multiple series from the same entity.

To this end, we propose a temporal adaptation of the MAML framework tailored to cold-start forecasting. Our approach introduces two key components: (i) temporal support–query splits within individual series, and (ii) temporal inner-loop updates during meta-training. Together, these yield a meta-learned initialization explicitly optimized for rapid adaptation to new demand series and efficient updating as their observed histories expand. We evaluate the proposed method on a diverse set of benchmark time series datasets and across multiple established deep forecasting architectures.

## References

- [1] Rob J Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018.
- [2] Marco Caserta and Luca D’Angelo. Intermittent demand forecasting for spare parts with little historical information. *Journal of the Operational Research Society*, 76(2):294–309, 2025.
- [3] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [4] Gang Wu, Li Cong, Chengbin Huang, Ying Ju, Jiange Jiang, and Chen Chen. Meta-learning framework for effective few shot time series prediction. In *2025 IEEE 5th International Conference on Power, Electronics and Computer Applications (ICPECA)*, pages 18–22. IEEE, 2025.

# Degradation-aware spare parts management using Deep Reinforcement Learning

Naim Al Khoury

Ghent University, Department of Industrial Systems Engineering and Product Design

e-mail: [naim.alkhoury@ugent.be](mailto:naim.alkhoury@ugent.be)

Dieter Claeys

Ghent University, Department of Industrial Systems Engineering and Product Design

e-mail: [dieter.claeys@ugent.be](mailto:dieter.claeys@ugent.be)

Zaharah Bukhsh

Eindhoven University of Technology,

Department of Industrial Engineering and Innovation Sciences

e-mail: [z.bukhsh@tue.nl](mailto:z.bukhsh@tue.nl)

Effective spare parts inventory management necessitates a delicate trade-off between maximizing service levels and minimizing ownership costs. This challenge is particularly sensitive in the service industry, where service providers maintain large fleets of equipment across a diverse customer base [1]. The complexity of this environment is compounded by the stochastic nature of replenishment lead times. Although real-time monitoring via Condition-Based Maintenance (CBM) offers valuable Advance Demand Information (ADI), traditional inventory policies generally do not incorporate this valuable information.

To the best of our knowledge, the Proactive Base Stock Policy (ProBSP) is the first degradation-aware policy to successfully integrate degradation data for complex systems involving multiple machines and stochastic lead times [2]. While the findings suggest that ProBSP effectively reduces holding costs by up to 67% compared to the Base Stock Policy (BSP), the lack of alternative degradation-aware policies in this context limits the scope for comparative benchmarking. Moreover, the ProBSP does not support batch ordering, a limitation that restricts its applicability in realistic scenarios where parts are frequently consolidated into batches.

To address these limitations, we learn spare parts policies by leveraging Deep Reinforcement Learning (DRL). DRL is a machine learning technique that has demonstrated considerable promise in the context of learning optimal policies for complex inventory problems, as evidenced by the relevant literature, including the works of [3] and [4]. We propose a comprehensive framework for developing and benchmarking degradation-aware inventory policies under realistic conditions, specifically incorporating batch ordering alongside multiple machines and stochastic lead times. This study, therefore, establishes a foundation for future research in the field of degradation-aware spare parts management. Within the proposed framework, we evaluate the performance of three state-of-the-art DRL algorithms, investigating their ability to unlock the full potential of degradation

data. Furthermore, we enrich these algorithms with domain knowledge, including insights from the BSP and the ProBSP.

We conduct an extensive numerical study to benchmark the performance of DRL-based policies against both the ProBSP and the modified BSP [5]. The experiments have been designed to replicate the heterogeneity of real-world systems. They span a wide range of operational settings, including varying numbers of machines, cost structures, and lead time distributions. Additionally, we emphasize the practical deployability of these methods by prioritizing algorithmic robustness over exhaustive hyperparameter tuning. Beyond benchmarking the economic benefits, we analyze the sensitivity of the policies to different model parameters and knowledge sources. Finally, we investigate the scalability of the proposed framework to large-scale problems and unravel the underlying decision mechanisms, thereby enhancing the explainability of the policies to facilitate their acceptance in industry.

Our results demonstrate that policies generated by the Deep Controlled Learning algorithm supported with knowledge from the ProBSP achieve cost savings of up to 25% against the modified Base Stock Policy and 15% against the ProBSP. These gains are driven by the effective integration of domain knowledge and the agent’s ability to dynamically exploit batch ordering opportunities. Furthermore, the framework proves highly scalable, maintaining consistent performance improvements for systems comprising up to 100 machines.

## References

- [1] Topan, E., A.S. Eruguz, W. Ma, M.C. van der Heijden, and R. Dekker (Apr. 2020). “A review of operational spare parts service logistics in service control towers”. en. In: *European Journal of Operational Research* 282.2, pp. 401–414. doi: 10.1016/j.ejor.2019.03.026.
- [2] Al Khoury, Naim, Dieter Fiems, El-Houssaine Aghezzaf, and Dieter Claeys (Apr.2024). “The Proactive Base Stock Policy for Condition Based Maintenance”. In: *Computers & Industrial Engineering* 190, p. 110077. issn: 0360-8352. doi: 10.1016/j.cie.2024.110077.
- [3] Gijsbrechts, Joren, Robert N. Boute, Jan A. Van Mieghem, and Dennis J. Zhang (May 2022). “Can Deep Reinforcement Learning Improve Inventory Management? Performance on Lost Sales, Dual-Sourcing, and Multi-Echelon Problems”. In: *Manufacturing & Service Operations Management* 24.3, pp. 1349– 1368. issn: 1526-5498. doi: 10.1287/msom.2021.1064.
- [4] Temizöz, Tarkan, Christina Imdahl, Remco Dijkman, Douniel Lamghari-Idrissi, and Willem van Jaarsveld (July 2025). “Deep Controlled Learning for Inventory Control”. In: *European Journal of Operational Research* 324.1, pp. 104–117. issn: 0377-2217. doi: 10.1016/j.ejor.2025.01.026.
- [5] Federgruen, A. and P. Zipkin (May 1986). “An Inventory Model with Limited Production Capacity and Uncertain Demands II. The Discounted-Cost Criterion”. In: *Mathematics of Operations Research* 11.2, pp. 208–215. issn: 1526-5471. doi: 10.1287/moor.11.2.208.

# Predictive Maintenance for Industrial Cinema Projectors Using Machine Learning

Subhamoy Dam

IESEG School of Management, Marketing and Sales

e-mail: [s.dam@ieseg.fr](mailto:s.dam@ieseg.fr)

Kristof Coussement

IESEG School of Management, Marketing and Sales

e-mail: [k.coussement@ieseg.fr](mailto:k.coussement@ieseg.fr)

For cinema projector manufacturers, equipment reliability is a critical competitive differentiator. Unplanned failures disrupt cinema operations, damage end-user experience, and generate significant costs through warranty claims and emergency field service. As projector systems grow increasingly sophisticated with laser-based technologies, the need for proactive maintenance strategies has become urgent. Despite widespread adoption of AI-driven predictive maintenance in manufacturing and energy sectors, applications in cinema projection systems remain unexplored. A review of the literature (Dalzochio et al., 2020) reveals that existing research focuses on industrial machinery, automotive systems, and aircraft engines, with few works comprehensively integrating traditional methods, machine learning, deep learning, multi-time horizon prediction, and feature importance analysis. Moreover, most predictive maintenance studies target failure prediction for individual components using a single measurement criterion—such as vibration, temperature, or rotational speed—through approaches like remaining useful life estimation, run-to-failure analysis, or condition monitoring. Research addressing system-level failure prediction that synthesizes heterogeneous sensor data across multiple subsystems remains limited.

This study, conducted in collaboration with a leading cinema projector manufacturer, develops a predictive maintenance framework for Laser Diode Module (LDM) failures. Unlike prior work focusing on single-component, single-sensor prediction, our approach predicts system-level failures by integrating heterogeneous sensor data from multiple subsystems. We benchmark traditional cross-sectional methods, ensemble machine learning approaches, and deep learning architectures, implementing multi-time horizon prediction and feature importance analysis. Using a subset of proprietary telemetry data comprising 557,653 observations across 130+ sensors spanning voltage, current, temperature, and rotational speed measurements, we implement a session identification strategy that distinguishes active and idle periods, aggregating readings into 24-hour runtime windows to reduce dimensionality and address class imbalance. The processed dataset yields over 24,000 sliding window instances, evaluated using  $5 \times 2$  cross-validation with hyperparameter tuning optimized for F1-score.

This research contributes empirical insights into algorithm selection for specialized equipment monitoring. As a first step, this study predicts LDM failures for a specific projector model. Future research will scale the dataset and expand to other models and failure types across the manufacturer's product portfolio.

## Keywords

Predictive Maintenance, Machine Learning, Cinema Projectors, Laser Diode Module, Time Series Classification, Deep Learning, Industrial IoT

## References

- [1] Dalzochio, J., Kunst, R., Pignaton, E., Binotto, A., Sanyal, S., Favilla, J., & Barbosa, J.. "Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges." *Computers in Industry*, 123, 103298. 2020.

# Determining extremal chemical graphs of maximum degree at most 3 via a polyhedral description

Valentin Dusollier

University of Mons,  
Computer Science Department - Algorithms Lab  
e-mail: [valentin.dusollier@umons.ac.be](mailto:valentin.dusollier@umons.ac.be)

Sébastien Bonte

University of Mons,  
Computer Science Department - Algorithms Lab  
e-mail: [sebastien.bonte@umons.ac.be](mailto:sebastien.bonte@umons.ac.be)

Gauvain Devillez

University of Mons,  
Computer Science Department - Algorithms Lab  
e-mail: [gauvain.devillez@umons.ac.be](mailto:gauvain.devillez@umons.ac.be)

Alain Hertz

Polytechnique Montréal - Gerad, Montréal, Canada  
Department of Mathematics and Industrial Engineering  
e-mail: [alain.hertz@gerad.ca](mailto:alain.hertz@gerad.ca)

Hadrien Mélot

University of Mons,  
Computer Science Department - Algorithms Lab  
e-mail: [hadrien.melot@umons.ac.be](mailto:hadrien.melot@umons.ac.be)

David Schindl

University of Applied Sciences Western Switzerland, Genève, Switzerland  
Haute Ecole de Gestion de Genève  
e-mail: [david.schindl@hesge.ch](mailto:david.schindl@hesge.ch)

A chemical graph is a simple, connected graph of maximum degree 3 or 4 (depending on the context). Such graphs represent molecules, where vertices are associated to atoms and edges to chemical bonds. Degree-based topological indices play a fundamental role in mathematical chemistry by capturing the structural properties of molecules and predicting their physicochemical behavior. These indices are calculated as the sum of the edge weights of the chemical graph, each edge having a weight defined by a formula that depends only on the

degrees of its endpoints. Numerous scientific papers have been published with the primary objective of characterizing chemical graphs that optimize (maximize or minimize) one of these degree-based topological indices. Restricting our attention to chemical graphs with maximum degree at most 3, we show that whatever the degree-based topological index (more than fifty have already been defined, and new ones are proposed regularly), it is sufficient to evaluate at most 16 chemical graphs to determine an optimal one. This follows from a complete polyhedral description of all chemical graphs with  $n$  vertices,  $m$  edges and maximum degree at most 3. Our main result is that whatever  $n$  and  $m$ , the associated polytope has at most 16 extreme points, and at least one of them corresponds to the structure of an optimal molecule.

# On the approximability of the balanced minimum evolution problem

Daniele Catanzaro

Université Catholique de Louvain, Belgium, CORE

e-mail: [daniele.catanzaro@uclouvain.be](mailto:daniele.catanzaro@uclouvain.be)

Raffaele Pesenti

Ca' Foscari Università di Venezia, School of Management

e-mail: [pesenti@unive.it](mailto:pesenti@unive.it)

Francesco Pisanu

CORE, Université Catholique de Louvain, Belgium

e-mail: [francesco.pisanu@uclouvain.be](mailto:francesco.pisanu@uclouvain.be)

## 1 Introduction

Consider a set  $\Gamma = \{1, 2, \dots, n\}$  of  $n \geq 3$  items and let  $\mathbf{D} = \{d_{ij}\}$  be a *dissimilarity* matrix on  $\Gamma$ , i.e., a symmetric matrix with all diagonal entries equal to 0 and nonnegative off-diagonal entries. An *Unrooted Binary Tree* (UBT)  $T$  on  $\Gamma$  is a tree having  $\Gamma$  as its leaf set and internal vertices of degree 3. Let  $\Theta_n$  denote the set of all UBTs on  $\Gamma$ , and for a given UBT  $T \in \Theta_n$ , let  $\tau_{ij}$  denote the number of edges on the unique path in  $T$  between items  $i$  and  $j$ . Then, the *Balanced Minimum Evolution Problem* (BMEP) seeks a UBT  $T_{\mathbf{D}}^* \in \Theta_n$  that solves the following nonlinear network design problem:

$$\min_{T \in \Theta_n} L_{\mathbf{D}}(T) = \sum_{i \in \Gamma} \sum_{\substack{j \in \Gamma \\ j \neq i}} \frac{d_{ij}}{2^{\tau_{ij}}}.$$

The BMEP arises from the literature on molecular phylogenetics [4]. In this setting,  $\Gamma$  represents a collection of distinct aligned molecular sequences (commonly referred to as *taxa*), such as DNA, RNA, codon sequences, or entire genomes;  $\mathbf{D}$  encodes estimated evolutionary distances between pairs of taxa [2]; and a UBT on  $\Gamma$  represents a candidate *phylogeny*, i.e., a tree hypothesizing the evolutionary relationships among the considered taxa.

From a complexity-theoretic perspective, Fiorini and Joret [3] proved that the BMEP is  $\mathcal{NP}$ -hard, via a reduction from the 3-coloring problem, and inapproximable within any factor  $c^n$ , for some constant  $c > 1$ , unless  $\mathcal{P} = \mathcal{NP}$ . These hardness and inapproximability results rely on reductions that construct instances of the BMEP in which  $\mathbf{D}$  contains zero-valued off-diagonal entries. Such cases, however, are rarely encountered in practical phylogenetic analyses, where

taxa are typically assumed to be all distinct (leading so to dissimilarity matrices having strictly positive off-diagonal entries).

## 2 Contributions

A dissimilarity matrix  $\mathbf{D}$  on  $\Gamma$  is *additive* if for every distinct  $i, j, p, q \in \Gamma$ , precisely one of the following *four point conditions* holds [1]:

$$\begin{aligned} d_{ij} + d_{pq} &\leq d_{ip} + d_{jq} = d_{iq} + d_{jp}, \\ d_{ip} + d_{jq} &\leq d_{ij} + d_{pq} = d_{iq} + d_{jp}, \\ d_{iq} + d_{jp} &\leq d_{ij} + d_{pq} = d_{ip} + d_{jq}. \end{aligned}$$

A tree  $T$  is *additive* for  $\mathbf{D}$  if:

$$d_{ij} = \sum_{e \in \text{path from } i \text{ to } j} w(e) \quad \forall \text{ distinct } i, j \in \Gamma.$$

Buneman [1] shows that the set of additive matrices on  $\Gamma$  and the set of additive trees on  $\Gamma$  are the same.

Thus, we answer to the natural question of whether this result could also serve as a certificate of optimality for additive instances of the BMEP.

**Theorem 1** *If  $\mathbf{D}$  is additive then any corresponding additive tree is the optimal solution of the BMEP.*

Independently, we then show that when all pairwise distances are strictly positive and confined within a prescribed interval, the problem admits a polynomial-time approximation algorithm, whose performance guarantee is explicitly linked to the width of this interval.

**Theorem 2** *If the off-diagonal entries of  $\mathbf{D}$  belongs to  $[a, b]$ , with  $0 < a < b$ , then the BMEP admits an approximation factor within  $b/a$ .*

Altogether, these results provide new insight into the approximability of the BMEP in settings that are directly relevant to phylogenetic inference.

## References

- [1] P. Buneman, “The recovery of trees from measures of dissimilarities,” in *Archaeological and Historical Science*, F. R. Hodson, D. G. Kendall, and P. Tautu (eds.), Edinburgh University Press, Edinburgh, UK, pp. 387–395, 1971.
- [2] J. Felsenstein, *Inferring Phylogenies*, Sinauer Associates, Sunderland, MA, 2004.
- [3] S. Fiorini and G. Joret, “Approximating the balanced minimum evolution problem,” *Operations Research Letters*, vol. 40, no. 1, pp. 31–35, January 2012.
- [4] O. Gascuel and M. A. Steel, *Reconstructing Evolution*, Oxford University Press, New York, NY, 2007.

# Genetic Branching: An Evolutionary Framework for Interpretable Branching Strategies

Simon Renard

Université libre de Bruxelles, Département d'informatique

e-mail: [simon.renard@ulb.be](mailto:simon.renard@ulb.be)

Quentin Louveaux

Université de Liège, Department of Electrical Engineering and Computer Science

e-mail: [q.louveaux@uliege.be](mailto:q.louveaux@uliege.be)

Bernard Fortz

HEC Liège, Université de Liège

e-mail: [bernard.fortz@uliege.be](mailto:bernard.fortz@uliege.be)

Solving mixed-integer linear programs (MILPs) plays a critical role in sectors ranging from energy planning to routing, with the Branch and Bound (B&B) algorithm serving as the primary method for exact resolution. The efficiency of B&B relies heavily on heuristics, particularly the branching strategy, which determines how the solution space is recursively divided. While Strong Branching (SB) is known to yield to small tree size, its computational intensity makes it impractical for full deployment, leading modern solvers like SCIP to rely on Reliability Branching (RB), a hybrid approach that combines SB with fast pseudocost estimates. Recently, the research community has investigated Machine Learning (ML) to design superior branching policies, utilizing methods such as Imitation Learning with Graph Convolutional Neural Networks (GCNN) or Reinforcement Learning. However, these approaches result in "black-box" models that lack interpretability and require large datasets, often thousands of instances, for training.

In this work, we introduce Genetic Branching (GB), a framework based on genetic programming to automatically design branching strategies that are both efficient and interpretable. Unlike neural network approaches, GB evolves symbolic scoring functions represented as abstract syntax trees (AST). These functions assign a score to each candidate variable, with branching performed on the variable maximizing this score. The evolutionary process is driven by a population of individuals that combine static, dynamic, and tree-based features using standard arithmetic and logical operators.

Our experimental results on four different sets of problems demonstrate that GB consistently matches or outperforms existing methods even on instances larger than those seen during training. In addition, it requires significantly fewer training instances and leads to interpretable functions.

# A New Algorithm to Enumerate the Vertices of the Balanced Minimum Evolution Polytope

Daniele Catanzaro

Université Catholique de Louvain  
Center for Operations Research and  
Econometrics  
daniele.catanzaro@uclouvain.be

Brieuc Pierre

Université Catholique de Louvain  
Center for Operations Research and  
Econometrics  
brieuc.pierre@uclouvain.be

Roberto Ronco

National Research Council of Italy  
Institute of Marine Engineering  
roberto.ronco@cnr.it

January 29, 2026

Consider a set  $\Gamma = \{1, 2, \dots, n\}$  of  $n \geq 3$  items. An *Unrooted Binary Tree* (UBT)  $T$  on  $\Gamma$  is a tree having  $\Gamma$  as its leaf set and internal vertices of degree 3. Let  $\Theta_n$  denote the set of all UBTs on  $\Gamma$ , and for a given UBT  $T \in \Theta_n$ , let  $\tau_{ij}$  denote the number of edges on the unique path in  $T$  between items  $i$  and  $j$ . Then, given a *dissimilarity* matrix  $\mathbf{D} = \{d_{ij}\}$  on  $\Gamma$ , the *Balanced Minimum Evolution Problem* (BMEP) seeks a UBT  $T_{\mathbf{D}}^* \in \Theta_n$  that solves the following nonlinear network design problem:

$$\min_{T \in \Theta_n} L_{\mathbf{D}}(T) = \sum_{i \in \Gamma} \sum_{\substack{j \in \Gamma \\ j \neq i}} \frac{d_{ij}}{2^{\tau_{ij}}}. \quad (1)$$

Recent advances on the polyhedral combinatorics of the problem [1] enabled the identification of fundamental characteristics of the convex hull of the solutions to the BMEP (also referred to as the *BMEP polytope*) as well as the description of some of its facets. These results have been possible also thanks to the assistance of general purpose softwares for polyhedral analyses, such as Polymake. The running time of these softwares, however, become quickly unpractical for increasing values of  $n$ . Hence, in this article, we address the problem of designing tailored algorithms for the study of the BMEP polytope, with special focus on the problem of enumerating its vertices. In particular, starting from the results discussed in [2], we present here a novel algorithm for enumerating distinct UBTs which proves significantly faster than the current state-of-the-art algorithm presented in [2]. This algorithm can be easily parallelized and, if appropriately modified, it can be converted into an implicit enumeration algorithm for the BMEP with performance competitive with the current state-of-the-art exact solver for the problem.

## References

- [1] D. Catanzaro, R. Pesenti, A. S. Barboza, and L. A. Wolsey. “Optimizing over Path-Length Matrices of Unrooted Binary Trees”. *LIDAM Discussion Paper CORE - 2023/20*, 2024.
- [2] D. Catanzaro and R. Pesenti. “Enumerating vertices of the balanced minimum evolution polytope”. *Computers & Operations Research*, vol. 109, pp. 209-217, 2019. doi: 10.1016/j.cor.2019.05.001.

## Index

- Aerts, M., 165  
Aghezzaf, E., 136, 192, 226  
Agnētis, A., 86  
Aguayo Quintana, C., 180  
Ahmadi, Z., 43  
Aktas Dejaegere, D., 154  
Al Khoury, N., 252  
Alsalam, A., 114  
Annaert, R., 190  
Arda, Y., 233, 235  
Arts, J., 211  
Aziz, H., 164
- Baecke, P., 142  
Baesens, B., 170, 205  
Bahadornia, M., 64  
Ballegeer, M., 71  
Barbini, A., 248  
Bayri, S., 245  
Bebronne, E., 182  
Becker, T., 239  
Benoit, D., 69, 71, 74, 144, 167  
Berghman, L., 132  
Berkman, B., 120
- Bessemans, P., 59  
Biro, P., 164  
Blaise, L., 148  
Blondeel, P., 80  
Bodur, M., 184  
Bogaert, M., 72, 76, 250  
Bonte, S., 256  
Boute, R., 130, 174  
Bouten, L., 106  
Braekers, K., 61, 64, 118, 245  
Brunieau, G., 146  
Bukhsh, Z., 252  
Bérard, R., 165
- Caljon, D., 41  
Campbell, A.M., 63  
Caris, A., 98, 245  
Castillo Reyes, G., 104  
Catanzaro, D., 258, 261  
Cedro, M., 203  
Chaudhary, S., 112  
Claeys, D., 252  
Cools, M., 182  
Coppé, V., 148, 243

Cornelissens, T., 64  
Coulter, B., 239  
Coussement, K., 106, 176, 254  
Creemers, S., 94, 152  
Csaji, G., 164

D'Haen, R., 63, 64, 118  
Dam, S., 254  
Danoy, G., 49  
Davari, M., 184  
De Backker, T., 130  
De Baets, B., 248  
De Bock, K.W., 72  
De Cleen, T., 142  
De Kerpel, L., 144  
De Landtsheer, R., 53, 128  
De Lange, S., 72  
De Latte, J., 247  
De Moor, J., 221  
De Smedt, J., 221, 223  
De Turck, K., 112  
De Vos, S., 205  
De Vuyst, S., 45, 228, 230  
De Walsche, N., 89  
De Weerd, J., 43, 217, 221, 223  
Debevere, T., 116  
Deconinck, G., 209  
Defryn, C., 82, 102, 195  
Delecluse, A., 55  
Deleye, P., 195  
Demeulemeester, T., 164  
Demirtas, S., 120  
Depaire, B., 237  
Depoortere, J., 173  
Deprez, B., 39, 165  
Derval, G., 47  
Devillez, G., 256  
Devos, A., 116  
Devriesere, K., 158, 160  
Diaz Gonzalez Manjarrez, F., 106  
Diels, J., 219  
Driesen, J., 173  
Dufour, M., 124  
Dumon, M., 174  
Dumont, M., 59, 114, 178

Dusollier, V., 256

Ehmke, J.F., 63  
El Habti, O., 118  
Elsherif, O., 98  
Ernst, D., 47  
Ernst, J., 114  
Evgeniou, T., 168  
Exoo, G., 109

Feuerriegel, S., 20, 41  
Fiems, D., 112  
Fontaine, G., 146  
Fontan, F., 124  
Forte, L., 228  
Fortz, B., 260  
François, V., 233, 235

Galarza Montenegro, B., 98, 188  
Garcia Tercero, L., 92  
Gendreau, M., 235  
Germeau, F., 53  
Ghorbani Saber, R., 226  
Gilet, F., 152  
Gillard, X., 53  
Glado, A., 78  
Goedertier, F., 142  
Goedgebeur, J., 109  
Goethals, A., 205  
Goethals, S., 168, 203  
Goossens, D., 84, 158, 160  
Goossens, L., 247  
Gürsoy, A., 134

Hadj Salem, K., 90  
Han, Y., 111  
Hanusa, R., 45  
He, X., 223  
He, Y., 201, 203  
Hedayati, S., 230  
Hertz, A., 256  
Hinns, J., 168, 203  
Hirsch, L., 78  
Hojny, C., 162

Ichmoukhamedov, T., 203  
Innocente, E., 122

Jans, M., 126  
Janssens, B., 67  
Janssens, W., 250  
Jiménez, D., 96  
Jooken, J., 109  
Joubert, J., 61, 100  
Jurgens, D., 170

Kameugne, R., 53, 150  
Kamhi, H., 132  
Karabag, O., 120  
Kazmi, H., 171, 173, 209  
Kerff, A., 178  
Keutgen, G., 146  
Khayyati, S., 120  
Kourentzes, N., 213, 215  
Köhler, C., 63

La Greca, A., 108  
Lahrs, L., 157  
Langenaeker, A., 102  
Lanzarone, E., 134  
Laporte, G., 20  
Lauwens, B., 80  
Legrand, E., 150  
Leloup, E., 138  
Lessmann, S., 205  
Leus, R., 84, 86, 225, 240  
Leyman, P., 226, 228  
Li, Y., 74  
Limbourg, S., 180, 182  
Limère, V., 136, 192  
Louveaux, Q., 47, 260

Maarouf, A., 41  
Maharani, A., 233  
Marchese, A., 78, 207  
Marckx, N., 76  
Martens, D., 168, 201, 203, 205  
Martens, J., 88  
Martin, N., 126, 237  
Marx, A., 247  
Mascarenhas, M., 171  
Maudet, G., 49  
Medina-Olivares, V., 205  
Mehmood, F., 209

Meire, M., 106  
Meurisse, Q., 53, 128  
Miftari, B., 47  
Moonen, C., 180  
Mouffe, M., 146  
Mues, C., 205  
Mélot, H., 256  
Münnich, R., 114

Nigudkar, S., 82  
Nimmegeers, P., 102

Ortega Vázquez, C., 165  
Oyamada, R., 223

Pagnoncelli, B., 96, 184  
Pan, C., 149  
Paquay, C., 59, 138  
Perneel, E., 86  
Pesenti, R., 258  
Pierre, B., 261  
Pinto, R., 134  
Pironet, T., 138  
Pisanu, F., 258  
Porbar, G., 136, 192  
Puchinger, J., 240

Qaiser, T., 209

Raa, B., 45, 230  
Rahman, S., 67  
Ramaekers, K., 64  
Rauh, F., 156  
Renard, S., 260  
Reusens, M., 170  
Reussens, Q., 247  
Riebus, M., 126  
Roa, J., 219  
Ronco, R., 261  
Roose, D., 104

Sagaert, Y., 215  
Salvadori, I., 86  
Schaus, P., 53, 55, 150  
Schindl, D., 256  
Seeholzer, L., 61  
Serral, E., 43, 219

Sitters, R., 51  
Smet, P., 134  
Spieksma, F., 162  
Stephan, T., 94  
Stubbe, L., 109  
Stützle, T., 197  
Sörensen, K., 98, 140, 194

Tan, B., 120  
Tancrez, J., 122, 152  
Tassin, L., 235  
Thanos, E., 136, 192  
Tharwat, H., 237  
Thuy, A., 144  
Tolleneare, J., 88

Van Belle, J., 41, 174, 213  
Van Bulck, D., 158, 160  
Van Camp, T., 69  
Van Den Abeele, M., 140  
Van Den Eede, T., 109  
Van den Poel, D., 72, 76, 250  
Van Der Auweraer, S., 176, 211  
Van Der Haar, J., 130, 174  
Van Der Veecken, S., 168  
Van Ee, M., 51  
Van Hentenryck, P., 55  
Van Hertem, D., 57  
Van Hoeck, I., 65  
Van Mulders, D., 76  
Van Pelt, T., 211  
Van Utterbeeck, F., 80  
Vanden Berghe, G., 104, 134, 239  
vanden Broucke, S., 217, 223

Vandendriessche, B., 167  
Vandomme, E., 178, 235  
Vansteenwegen, P., 61, 65, 154, 241  
Verbeke, W., 39, 41, 205, 213  
Verboven, S., 78, 207  
Verdonck, T., 39, 205  
Verplancke, H., 136, 192  
Verstraete, J., 241  
Verstraeten, A., 171  
Verwaeren, J., 248

Walraevens, J., 116  
Wang, H., 84  
Wang, Y., 186  
Wauters, T., 90, 108  
Wei, W., 199  
Wens, M., 241  
Wessel, S., 162  
Weymeirsch, J., 114, 178  
Weytjens, H., 221  
Wittevrongel, S., 116  
Wolfe, A., 176  
Wuyts, B., 217

Yaman, H., 57, 96, 111  
Yang, F., 84  
Yang, K., 225  
Yi, X., 84

Zhang, S., 199  
Zhao, B., 240  
Zong, S., 100

Çalik, H., 57, 239