

Unlocking OR for all: From models to business impact

6th EURO Practitioners' Forum Conference

Monday 27 April 2026 (12:30pm) - Tuesday 28 April 2026 (4:30pm)

Warsaw University of Technology

Organisation committee:

- Vladimir Fux (Zalando)
- Michele Quattrone (AirLiquide)
- Sander van Aken (Flix SE)
- Torkel Haufmann (Sintef)
- Waldemar Kocjan (Jeppesen ForeFlight)
- Izabela Żółtowska (Warsaw University of Technology)
- Mariusz Kaleta (Warsaw University of Technology)
- Susanne Heipcke (FICO) - chair



Contents

Programme	3
Sponsors	5
Keynote talks	6
Ciodyk (ORTEC): Practical optimization: real-world success stories	6
De Smet (Timefold): How to build an open source solver	6
Kłos (PSE): Framing fairness: Beyond-cost optimization for TSOs	7
Lichocki (Google): An engineering perspective on mixed integer programming	7
Stopa (IBM): Quantum computer – from holy grail to your favourite toy	8
Abstracts of contributed talks	9
Bauerhenne (ZDF): Data-driven TV scheduling	9
Ceyhan (Just Eat Takeaway): Interactive multi-objective optimization	9
Guerrero (IKEA): The human factor in OR	10
Held (Swissport): Optimizing vacation planning at Swissport	11
Balbiyad (Kpler): Unveiling maritime market opacity	11
Broihan (GAMS): The messy truth about deploying models	12
Benoist (Hexaly)/Quattrone (AirLiquide): Hexaly Studio	13
Kurzejamski (IDEAS): Municipal tax automation	14
Souza (Algomia): A flow-based solver for large-scale combinatorial optimization	14
Wojdan/Ziółkowski (Orange): Telco AI Assistant	15
Van Aken (Flix SE): Unlocking network-wide profitability steering	16
Fonseca (Univ. Coimbra) Computational modelling of a carrier assignment problem	17
Wiese (TUM): StudyPlanner: Helping students to plan university courses	18
Guajardo (NHH): Developing an optimization tool for farm advisory	18
Bacher (Zalando): Designing online optimisation systems	19
Wolff: Building an optimisation platform for everyone	19
Jurado/Wolf: Delivering optimization to non-experts	20
Pozzi (Optit): Research competitions as a catalyst for open innovation	21
Domański (Betacom): Internal transport optimization opportunities for warehouses	22
Gupta (CW): Digital twins for simulation and optimization	22
Panel discussion	24
Discussion group topics	25
OR + GenAI – how they complement each other	25
Getting solutions into production	25
Academia vs. Industry: The relevancy gap	25
Maps	26

Programme – *DRAFT*

Monday 27 April, 2026

10:30 Guided campus walking tour – group 1 (optional)

11:15 Guided campus walking tour – group 2 (optional)

12:00 Registration

12:30 Welcome session

12:45 **Keynote 1:** S. Ciodyk (ORTEC): *Practical optimization: real-world success stories*

13:35 Contributed talks (C1)

- C. Bauerhenne (ZDF): *Data-driven TV scheduling: Broadcasting ZDFinfo documentaries*
- G. Ceyhan (Just Eat): *Use of LLM-based agents for interactive multi-objective optimization*
- V. Guerrero (IKEA): *The human factor in OR: Strategies for building trust and effective onboarding in optimization projects*
- M. Held (Swissport): *Optimizing vacation planning at Swissport: Bridging technical solutions and organizational change*

14:55 Coffee break

15:20 **Keynote 2:** G. De Smet (Timefold): *How to build an open source solver and a company around it*

16:10 Contributed talks (C2)

- S. Balbiyad (Kpler): *Unveiling maritime market opacity: An Operations Research pipeline for product and trade reconstruction*
- J. Broihan (GAMS): *Confessions of an OR consultant: The messy truth about deploying models*
- T. Benoist (Hexaly)/M. Quattrone (AirLiquide): *Hexaly Studio: a platform for optimization prototypes and applications*
- G. Kurzejamski (IDEAS Research Institute): *Disrupting public sector IT: An agile, hardware-constrained approach to deploying local LLMs for municipal tax automation*

17:30 Parallel discussion groups

19:00 Close

20:00 Conference dinner (optional)

Tuesday 28 April, 2026

8:00 *Registration & coffee*

8:30 **Keynote 3:** M. Kłos (PSE): *Framing fairness: A practical approach to beyond-cost optimization for TSOs' mathematical programming problems*

9:20 Contributed talks (C3) – short presentations

- A. Souza (Algomia): *A flow-based solver for large-scale combinatorial optimization*
- I. Wojdan/M. Ziółkowski (Orange): *Telco AI assistant: From research to delivery – challenges and lessons learned*
- S. Van Aken (Flix SE): *Unlocking network-wide profitability steering – the case of automated network initialization*
- C.M. Fonseca (Univ. Coimbra): *Computational modelling of a carrier assignment problem under the ROAR-NET API*

10:00 Panel discussion (main lecture hall) | Talent clinic (breakout rooms)

10:40 *Coffee break*

11:10 **Keynote 4:** P. Lichocki (Google): *An engineering perspective on mixed integer programming in computer systems*

12:00 Contributed talks (C4)

- A. Wiese (TU Munich): *StudyPlanner: Helping students to plan university courses with integer linear programming*
- M. Guajardo (Norwegian School of Economics): *Developing an optimization tool for farm advisory*
- Ch. Bacher (Zalando): *"Believe me, it is optimal": Designing online optimisation systems for explainability & acceptance – a case study*

13:00 *Lunch*

14:20 **Keynote 5:** T. Stopa (IBM): *Quantum computer – from holy grail to your favourite toy*

15:10 Contributed talks (C5) – short presentations

- J.K. Wolff: *Building an Optimisation platform for everyone*
- M. Jurado/P. Wolf: *Delivering optimization to non-experts with structured dialogue, dataset profiling, and solver verification*
- M. Pozzi (Optit): *Bridging the gap in academia-industry collaboration: Research competitions as a catalyst for open innovation*
- P. Domański (Betacom): *Internal transport optimization opportunities for warehouses*
- R. Gupta (Computacionele Wetenschapsgroep): *Digital twins for simulation & optimization of heavy equipment multipart logistics*

16:00 Reports from discussion groups and closing remarks

16:30 *End of event*

18:00 *Guided tour of the old town center of Warsaw (optional)*

Sponsors

Many thanks to all the sponsors and institutional supporters of this event that has been entirely financed through their generous contributions.

Gold



Silver



Bronze



Institutional supporters



Keynote talks

Practical optimization: real-world success stories

Sylwester Ciodyk

ORTEC Poland

In today's fast-moving world, optimizing efficiency and effectiveness is essential for every company. Drawing on my experience designing and deploying optimization systems at ORTEC, I'll share field-tested stories of how we connected algorithms to human goals and operational constraints. Join me as we reveal how optimization drives business transformation, with real-life examples such as smarter load building in e-commerce and improved forecasting in the energy sector, leading to significant operational savings. Expect pragmatic take-aways you can apply immediately—how to quantify impact, handle uncertainty, and bridge the gap from prototype to production—to deliver measurable savings and better service.



With 15 years in Product and Project Management, **Sylwester Ciodyk** specializes in optimizing processes and driving efficiency across Automotive, Location Intelligence, Energy, and Retail. Currently, supporting the Director of Operations for the European Energy sector, Sylwester oversees key performance indicators, assists executive leadership, and leads initiatives focused on cost-effective solutions and automation. Known for combining strategic insight with practical execution, Sylwester helps organizations achieve measurable efficiency and growth.

How to build an open source solver and a company around it

Geoffrey De Smet

Timefold

<https://www.linkedin.com/in/geoffrey/>

Almost 20 years ago, I created an open source solver, now called Timefold Solver. These days, it is used across the globe to optimize vehicle routing problems, job shop scheduling, shift rostering and many other planning cases, directly or indirectly affecting the lives of millions of people.

In this keynote, I'll cover the challenges, successes and lessons learnt. What makes our solver unique. How our engineering culture handles trade-off decisions. And why the timing is right for all of us to make Operations Research mainstream.



Geoffrey De Smet is founder/CTO of Timefold and the creator of OptaPlanner, its predecessor. He's an expert in planning optimization, an active open source contributor and an international speaker. His technology is in use across the globe, to automate complex scheduling and routing of hundreds of thousands of employees, vehicles and machinery every day. Before founding Timefold, he worked as a Senior Principal Software Engineer at Red Hat (IBM), project lead at Schaubroeck and as an academic researcher at Kaho Sint-Lieven (KU Leuven).

Framing fairness: A practical approach to beyond-cost optimization for TSOs' mathematical programming problems

Michał Kłos

PSE

While Operational Research (OR) traditionally focuses on cost optimization, real-world applications often demand more than a single dimension of the objective function. In many real complex systems, the concept of "fairness" emerges as a crucial, yet subjective, element of the objective function. This talk explores how OR can be used to model and achieve equitable outcomes, even when "fairness" lacks a universal definition and must be translated into quantifiable objectives. We argue that incorporating fairness requires a deliberate shift from purely technical optimization toward a more nuanced approach that integrates stakeholders' perceptions to equity considerations.

We present two distinct case studies from the energy sector that illustrate this challenge: Cost Distribution of Congestion Income and Cost Sharing of Remedial Actions. Through these examples, we demonstrate that achieving a "fair" outcome is not an add-on but an integral part of the modeling process. We conclude that successful applications of mathematical programming in this domain require not only technical expertise but also a deep understanding of the human and policy factors that define what is perceived as a just and equitable solution.

Michał Kłos is the Deputy Director for Computational Method Development at the NCAE Bureau (Polskie Sieci Elektroenergetyczne S.A., the Polish TSO) and former Director of the Interdisciplinary Center for Energy Analyses (ICAE) in the National Centre for Nuclear Research. He specializes in mathematical modeling and optimization for energy markets and power systems, with a focus on implementing EU regulations.

His work involves developing practical algorithms, including the clustering algorithm used in the first ENTSO-E Bidding Zone Review. He also co-authored the IP-DSR service concept, being responsible for the contracting and operational dispatch algorithms for service providers. His expertise includes leading energy market modeling projects for Transmission System Operators and acting as an expert on cost-sharing methodologies and congestion income distribution.

A graduate of Computer Physics from Jagiellonian University, he has authored over a dozen scientific publications and is a regular speaker at the annual European Energy Markets conference.



An engineering perspective on mixed integer programming in computer systems

Paweł Lichocki

Google

In this talk, we present Mixed Integer Programming (MIP) from an informal, practitioner's perspective. MIP is a powerful, declarative approach to solving combinatorial optimization problems. One expresses the problem at hand via a mathematical model, while a "black-box" solver does the number crunching and finds a solution. In practice, building a MIP-based application entails a few challenges. The optimization model must faithfully capture the reality of the problem while remaining "solvable" by the available solvers. Moreover, the optimization component must be integrated with the production system in a clear and robust way. We describe these issues in examples distilled from real-life applications of MIP to large-scale computer systems. We discuss the art of crafting efficient formulations for typical requirements observed in computer systems, like fault tolerance or churn control. We provide practical hints on how to encapsulate and verify mathematical optimization within a larger production environment. Finally, we discuss the relevant engineering challenges such as input/output validation, error resilience, and the ease of adding future extensions. Overall, we hope to shed light on the sometimes overlooked practical aspects of applying principled mathematical optimization to live and large-scale systems.

Paweł Lichocki is a software engineer at Google in the Operations Research team. For the past +10 years, he has been contributing to dozens of projects relying on combinatorial optimization. In particular, he applied Mixed Integer Programming in many real-life, large-scale systems and products. His interests include both practical and theoretical aspects of optimization, as well as evolutionary biology, metaheuristics, and machine learning. He received a PhD from the School of Computer and Communication Sciences at EPFL in 2013 for work on the evolution of division of labor in multi-agent systems. Prior, he was a researcher in the Supercomputer and Networking Center in Poznań where he worked on parallel and distributed processing algorithms. Recently, he has been exploring the possibility of using machine learning to improve MIP solvers.



Quantum computer – from holy grail to your favourite toy

Tomasz Stopa

IBM, Poland

For over 40 years the universal quantum computer was a holy grail for scientists. Today, we already have noisy, intermediate scale quantum computers (NISQ) to play with. These current devices are already believed to allow getting results beyond reach for classical supercomputers.

In the talk we'll explore the current state of quantum computing technology on the example of IBM, we'll see how to play with these toys and why they should become your favourite ones starting tomorrow. Finally, we'll cover when the quantum crusaders plan to get the holy grail and how we believe it looks like.



Tomasz Stopa works as software development manager in IBM's Kraków Lab, Poland. He obtained his Ph.D. in theoretical solid state physics from AGH University of Science and Technology. IBM Master Inventor with 35 patents and scientific publications. As IBM Quantum Ambassador promotes quantum computing within industry, academia and high schools. Co-organizer of Kraków Quantum Computing Seminar (KQIS) and mentor in quantum computing hackathons.

Abstracts of contributed talks

Data-driven TV scheduling: Broadcasting ZDFinfo documentaries

Dr. Carolin Bauerhenne¹, Andrea Reckenthäler¹, Dr. Xenija Neufeld²

¹ZDF (Zweites Deutsches Fernsehen), Germany
bauerhenne.c@zdf.de

²Accso – Accelerated Solutions GmbH, Germany

Public-service media play a crucial role in informing, educating, and entertaining audiences across Europe. As a German broadcaster, ZDFinfo combines editorial expertise with advanced analytics to optimize TV scheduling. This talk presents the practical end-to-end implementation of a mathematical programming approach for scheduling – from long-term planning to single-title placement – sustainably integrated into editorial workflows. Using several years of audience data, we apply constraint programming to maximize audience reach while ensuring smooth transitions between documentaries and meeting key constraints: public-service obligations, compliance requirements such as youth protection, and market-driven aspects like prime-time switching behavior. The approach accounts for uncertainty in audience predictions for topics with sparse or volatile data and incorporates strategies to prevent over-concentration on historically successful content while fostering novelty and diversity. Finally, we share lessons on leveraging interactive analytics tools to support TV schedulers – enabling them to combine data-driven recommendations with editorial expertise to navigate the complexities of stakeholder-driven media environments.

Biography

Carolin Bauerhenne holds a PhD in Operations Research and is attending EURO for the third time. At ZDF, she manages various AI projects – including the one presented today.

Andrea Reckenthäler is Channel Manager at ZDFinfo, Germany's leading documentary TV channel, and oversees content syndication and planning across four public broadcasters in Germany, Austria, and Switzerland. She holds degrees in communication research, television journalism, and history.

Xenija Neufeld is an AI Engineer at Accso, where she develops AI solutions for clients facing complex challenges. She holds a PhD in Computer Science, specializing in long-term planning and reactive execution and brings experience from the video games industry.

Using LLM-based agents for interactive multi-objective optimization

Gökhan Ceyhan

Just Eat Takeaway, Berlin, Germany
gokhanceyhan01@gmail.com

At Just Eat Takeaway (JET), we solve many optimization problems in logistics that involve multiple, often conflicting criteria. For example, our real-time capacity management algorithm adjusts partners' delivery areas in response to anticipated demand surges to maintain service quality while serving as many orders as possible. Another example is courier-order assignment, which assigns incoming orders to couriers

and determines each courier's delivery sequence. This system optimizes a cost function that combines objectives such as travel-time minimization, delay minimization, and operating cost.

A common approach to solve multi-objective optimization problems in practice is to reduce the multiple objectives into a single scalar objective using a predetermined scalarization function, such as weighted-sums or penalty terms. In some cases, the problem owner specifies a strict order of objective importance, and then we solve the problem lexicographically. After the proof-of-concept phase, as the initial algorithms evolve into mature products, we may expose the key configuration parameters of such scalarizations to decision makers, expecting them to adjust these parameters as preferences or system behavior change. This approach has several drawbacks. First of all, a predetermined scalarization may fail to represent the full set of nondominated solutions (i.e., solutions for which no other solution is at least as good in every criterion and strictly better in at least one). Secondly, users may struggle to configure algorithms correctly because they may not fully understand how those parameters affect the outcomes. And, last but not least, searching over many parameters to find a satisfactory solution can be cumbersome or computationally expensive, making fine-tuning inefficient and limiting the value extracted from OR.

To address these challenges and enable more effective use of OR solutions by business stakeholders, we explore LLM (large language model)-based agents as a communication layer that translates user preferences into algorithm parameters and supports an interactive solution process for the multi-objective problem at hand. We develop expert agents that interact with users in natural language and can call the functions provided by the optimization system. Our goal is to help users better understand the solution space of their problems and to reduce iterations, time, and effort required to obtain a satisfactory solution. To evaluate our approach, we compare its performance with well-known methods in interactive multi-objective optimization.

Biography

Dr. Gökhan Ceyhan is a Principal Data Scientist at Just Eat Takeaway. He earned his Ph.D. in Industrial Engineering from Middle East Technical University in Ankara, Turkey. With over 10 years of experience as an OR practitioner, Dr. Ceyhan has applied his expertise across diverse industries, including energy markets, aviation, and on-demand delivery. His primary research interests include multi-objective optimization, mixed-integer programming, and recently, reinforcement learning.

The human factor in OR: Strategies for building trust and effective onboarding in optimization projects

Victoria Guerrero, Arturo Pérez Rivera

IKEA

Victoria.guerrero@ingka.ikea.com, Arturo.perez.rivera@ingka.ikea.com

The true success of an optimization product in a real world setting depends not just on its technical accuracy or robust infrastructure. Crucially, it relies on gaining stakeholder trust, understanding, and acceptance. Our presentation will address this vital gap, demonstrating how we bridge the divide between advanced technical solutions and their practical adoption, especially when new systems replace older, established ones.

We will explore the challenges and strategies involved in engaging stakeholders effectively throughout an OR project's lifecycle, from initial problem definition to daily operations. While stakeholder experience is invaluable, some requirements might stem from past limitations that can now be overcome. We'll discuss how to facilitate open dialogue and build trust to distinguish between essential requirements and flexible goals, and how to manage the natural risk aversion that often accompanies transitions from legacy systems. This includes outlining strategies for a well-planned, safe deployment and rollout.

Using our experience at IKEA, where our team developed an optimization product for Customer Order Allocation that is gradually replacing a rule based legacy solver, we will share practical mechanisms for building trust. This includes insights into effective testing, thorough analysis, and benchmarking strategies. Our goal is to provide actionable approaches that make advanced Operations Research solutions more accessible, explainable, and ultimately, more impactful for a broader, non-expert audience.

Biography

Victoria Guerrero is a Senior Data Scientist at Ingka Group (IKEA), specializing in operations research and data-driven decision making. She began her career working on the integration of renewable energy into electric power systems and, over the past six years, she has focused on improving logistics processes through analytics and optimization. Beyond her technical work, Victoria is passionate about the intersection of business, product, and decision science. She enjoys engaging with different stakeholders to increase adoption and make onboarding easier, while translating real needs into meaningful product improvements.

Optimizing vacation planning at Swissport: Bridging technical solutions and organizational change

Moritz Held¹, Peter Fusek², Roman Berner¹, Andreas Klinkert²

¹Swissport International Ltd., Switzerland

moritz.held@swissport.com, roman.berner@swissport.com

²ZHAW School of Engineering, Switzerland

peter.fusek@zhaw.ch, andreas.klinkert@zhaw.ch

As the world's largest ground handler, Swissport manages vacation planning within a highly dynamic and operationally demanding environment. Especially at its biggest site at Zurich Airport, a handful of planners are responsible for the distribution of the vacation demands of around 2,500 employees. With substantial leave entitlements, year-round operations and peak vacation demand coinciding with the busiest travel season, the planning process requires considerable coordination.

This paper examines the development and deployment of a Mixed Integer Linear Program (MILP) in cooperation with the Zurich University of Applied Sciences (ZHAW) to support vacation planning automation at Zurich. The focus lies on both the successes and challenges in bridging the gap between technical academic expertise and practical application.

The project originated from a bachelor thesis exploring the possibility of a MILP solution for the vacation planning process. Close collaboration with Swissport planners proved essential in translating complex business requirements into mathematical constraints. This co-creation process proved essential, as it allowed us to identify existing unwritten rules and procedures, ultimately shedding light on the human aspect of decision making of individual planners. Rather than simply automating existing workflows, the optimization exercise became a catalyst for organizational change, revealing opportunities to untangle unnecessarily complex decision-making processes.

A key innovation in our approach was the strong emphasis on transparency in presenting optimization results. Displaying results transparently builds trust in the model's logic while providing planners with the objective justification for difficult decisions. The impact was significant: reduced planning time and notably improved issue resolution.

Although key stakeholders were supportive, a number of users showed reservations toward the optimization outputs. This highlights the essential role of structured change management in operations research (OR) project implementation.

For OR practitioners, we recommend co-creating model constraints with receptive users from the outset, actively challenging existing procedures, and prioritizing solution explainability. Most importantly, practitioners must recognize that organizational change is sometimes a prerequisite for optimization success, not merely an afterthought. By treating optimization as a tool for both technical improvement and organizational evolution, practitioners can increase the accessibility and impact of OR solutions, ensuring sophisticated models translate into user adoption and genuine business value.

Biography

Moritz Held is an Optimization Specialist at Swissport in the Business Engineering and Project Management Team. In his 1.5 years at Swissport, he was responsible for the conceptualization, development and deployment of several planning- and optimization systems. Before that he graduated from ZHAW School of Engineering with a bachelor's degree in business administration and engineering, where he wrote his bachelor thesis on the automation of the vacation planning process.

Unveiling maritime market opacity: An Operations Research pipeline for product and trade reconstruction

Saad Balbiyad, Daniel Chemla, Leo Dureuil, Samir Naim, Paul Renart

Kpler, France

{ sbalbiyad | dchemla | ldureuil | snaim | prenaud }@kpler.com

In the realm of commodity market intelligence, the opacity of maritime logistics presents a significant challenge for market players seeking to understand global supply and demand dynamics. While Automatic Identification System (AIS) data provides raw vessel positioning, it lacks the granular cargo tracking details necessary to identify the specific product and quantity being transported and, eventually, the trades formed between different market players.

Kpler's clients—primarily interested in global flows within major maritime markets like Oil, Gas, and Dry Bulk—rely on bridging this gap to see the full picture. To reconstruct these global flows, Kpler employs a bottom-up methodology that begins at the single vessel level and

aggregates upwards. This process utilizes advanced Operations Research (OR) techniques to solve specific challenges at each stage of the reconstruction:

- Quantity estimation and voyage construction via Mixed-Integer Linear Programming (MILP): using vessel draft and vessel specification, the model estimates the quantities exchanged at each port
- Product identification via exhaustive solution enumeration using Constraint Programming (CP), followed by solution merging to compute the final product composition across all ports visited by the vessel
- Trade reconstruction: To infer the final buyer-seller relationships, the problem is mapped onto a specific graph structure and solved as a Minimum Cost Flow problem. This optimization finds the most logical path of commodities through the network, effectively reconstructing the trades that underpin the physical movements.

This sequential reconstruction pipeline is triggered by real-time events (e.g., new AIS positions, port reports, or contractual data like vessel chartering). Although this modular architecture ensures computational tractability and clear separation of responsibilities, we discuss the inherent challenges regarding coupling constraints—specifically, the difficulty of enforcing product-dependent directionality (import/export) during the initial quantity estimation phase. In production, the pipeline may be executed several hundred thousand times per day, which requires robust, modern tooling to orchestrate solver components, manage concurrency, monitor solution quality, and ensure system availability.

Biography

Daniel Chemla is a Director of Engineering at Kpler, working on algorithmic and optimization-based solutions for large-scale industrial systems. He holds a PhD in Operations Research from École des Ponts ParisTech (ENPC) and Université Paris 13 (2009–2012), focused on transportation and energy applications. His doctoral research addressed the optimization of shared bike systems, and he participated in ROADEF challenges on large-scale maintenance planning for nuclear fleets under thousands of scenarios. After his PhD, he worked on data-driven algorithmic systems for mobility analysis. Since joining Kpler in 2018, he has contributed to solver-based methods for product identification in maritime flows, system architecture, and the creation of a new Power business unit combining optimization and machine learning.

Saad Balbiyad is a senior Operations Research engineer at Kpler. Over the past seven years, spanning his time at Kpler and previously as a research engineer at EDF, he has contributed to multiple projects relying on operations research and optimization. His experience covers the energy sector—such as hydropower optimization, energy management systems (EMS), and nuclear outage planning—as well as maritime intelligence, including the reconstruction of global maritime commodity flows using constraint-based models. Throughout his work, he has applied a broad range of techniques, including metaheuristics, mixed-integer linear programming (MILP), column generation, constraint programming, stochastic dynamic programming, and graph theory.

Confessions of an OR consultant: The messy truth about deploying models

Justine Broihan

GAMS Software GmbH, Augustinusstr. 11b, 50226 Frechen, Germany
jbroihan@gams.com

"The model solved the problem to optimality, but the client never used it."

This is the nightmare scenario for every Operations Research practitioner. In this session, I will move beyond the theoretical elegance of algorithms to discuss the gritty reality of working as an OR consultant. This is a candid, "no sugar-coating" look at the end-to-end journey of delivering projects—from the messy ambiguity of initial specifications to the human hurdles of final deployment.

Drawing on experience from different projects, I will share stories of both hard-won successes and instructive failures. I will outline practical strategies for managing stakeholder expectations and ensuring model explainability. The goal is to provide awareness for common pitfalls in OR projects and how to maintain clarity and alignment throughout the project lifecycle, ensuring the final delivery is not just accurate, but adopted.

Biography

As a Senior Projects and Consulting Specialist at GAMS Software GmbH, **Justine Broihan** has spent the last three years navigating the end-to-end delivery of numerous Operations Research projects. Her role focuses on guiding clients to not only enhance their GAMS skills but to design and deploy robust optimization solutions that work in the real world. This practical consulting experience is built upon a strong academic foundation, comprising six years of research and teaching and a PhD in Operations Research.

Hexaly Studio: a platform for optimization prototypes and applications

Thierry Benoist¹, Michele Quattrone²

¹Hexaly, Paris, France.
tbenoist@hexaly.com

²AirLiquide R&D, Paris, France.
michele.quattrone@airliquide.com

Hexaly Optimizer is a global optimization solver that combines exact and heuristic methods to find near-optimal solutions in minutes. Hexaly Studio (<https://studio.hexaly.com>) is a web application released in 2023 and built on top of Hexaly Optimizer. It includes a code editor for writing and debugging optimization models, as well as a graphical interface for visualizing solutions. The optimization is done remotely on dedicated servers with Hexaly Optimizer. In this talk, we will give a demo of Hexaly Studio and show how it can be used in industry or for teaching to develop quick prototypes and deploy them in production in minutes.

Hexaly Studio

Hexaly Studio is an online code editor that combines a dynamic programming language with modeling features. The studio comes with a collection of models for classical operations research problems and a generator for complex variants of routing and scheduling problems. A Hexaly Studio user can write their own model or modify an example to connect their data, add custom constraints, and objectives. The optimization is performed remotely on a dedicated server by Hexaly Optimizer using a combination of exact techniques and heuristic approaches. Finally, the solution is transferred back to Hexaly Studio and can be visualized through graphical widgets in a dashboard. Several widgets are available for route visualization, activity planning, and displaying indicators, among other purposes. Once the prototype phase is complete, the mathematical model can be directly imported into an application through Hexaly Optimizer's API or accessed directly via a REST API call.

Hexaly Studio x Air Liquide: A semi-industrialized tactical tool for home healthcare logistics

Air Liquide R&D presents a tactical optimization tool designed for internal homecare delivery logistics—a complex, country-specific blend of "milk round" and spot activities serving over 2.1 million patients in their homes. Managing this activity is a significant operational challenge: it requires strictly respecting patient visit plans scheduled at predetermined frequencies (ranging from daily to monthly), while coping with an environment characterized by fluctuating patient density and shifts in medical prescriptions.

Developed using Hexaly Studio, the tool provides our small-but-skilled community of logistics coaches with a platform to evaluate existing route structures and assess the impact of territory adjustments or organizational changes. This technical choice served as a pragmatic bridge to transform an R&D model into a semi-industrialized solution, without the need for the investment typically required to develop a custom application from scratch.

While straightforward, the interface allows users to iteratively refine the targeted scenario by adjusting soft and hard constraints—such as overtime quotas, vehicle capacities, and route districting—to ensure that mathematical efficiency remains aligned with operational common sense. For instance, to minimize implementation friction, users can balance operational efficiency against the need to keep patient delivery days as stable as possible. This approach has already been successfully deployed across several subsidiaries, serving as both a guide and a catalyst for data-driven decisions.

Biography

Before launching Hexaly, **Thierry Benoist** headed the Operations Research (OR) team at Bouygues, one of the biggest French public corporations. Graduated from École Polytechnique, he holds a Ph.D. in computer science from Avignon University (2004) and a Habilitation from Nantes University (2014). He has published papers in top OR journals. Several awards distinguished his research: 2005 Doctoral Prize for Innovative Applications by the French Computing Society, 2006 Robert Faure 3rd Prize by the French Operations Research Society (ROADEF), and finalist of the 2012 EURO Excellence in Practice Award.

Michele Quattrone is an OR Engineer within the Industrial Performance group at Air Liquide R&D. For twenty years, he has worked on product distribution logistics across the Group's diverse supply chains—spanning industrial cylinders, bulk liquids, and energy. His focus remains on bridging the gap between mathematical modeling and the reality of industrial operations. Michele is also a member of the organizing committee for the EURO Practitioners' Forum.

Disrupting public sector IT: An agile, hardware-constrained approach to deploying local LLMs for municipal tax automation

Grzegorz Kurzejamski

IDEAS Research Institute, Poland
grzegorz.kurzejamski@ideas.edu.pl

Government IT projects are notorious for their "waterfall" inefficiency: long centralized planning often result in obsolete systems that fail to meet user needs. This presentation demonstrates a radical departure from this norm: the "ExtrAct" project, implemented by **IDEAS Research Institute**, applies a "startup-style" methodology to public administration. We address the classic Operations Research challenge of **resource constraint optimization** by focusing on minimizing the workload of two critical human agents: municipal clerks and AI engineers. First, to optimize the engineering effort, we replaced manual prompt tuning with a **semi-automatic optimization loop**. In this system, massive "thinking" models analyze extraction errors and iteratively refine a static **query tree graph** for the smaller local models running on consumer-grade Edge hardware. This automation drastically reduces the software engineering time required to handle complex legal syntax. Second, we optimized the civil servants' workflow by shifting their role from manual data entry to rapid validation, effectively removing the process bottleneck. This dual optimization approach enabled a small team to deploy a high-value, privacy-compliant solution in just a few months, with minimal costs and a lean operational footprint, effectively becoming a reference point for our subsequent rapid Proof of Concept initiatives in the public sector.

Biography

Grzegorz Kurzejamski is an R&D Team Leader at IDEAS Research Institute and an AI Solutions Architect with over a decade of experience bridging academia and industry. He has authored multiple research grants (NCBR, PARP) and founded tech startups. A graduate of the Warsaw University of Technology, focuses on deploying cutting-edge AI in practical environments.

A flow-based solver for large-scale combinatorial optimization

Alexander Souza, Tristan Koning, Ori Chen

Algomia GmbH, Schachenallee 29, CH-5000 Aarau
{ alexander | tristan | ori }@algomia.com

This abstract reports our current engineering progress in designing, implementing, and operating a flow-based solver for large-scale and real-world combinatorial optimization. The solver has been tested extensively at the Swiss Federal Railways (and other customers). The challenge there was to optimize the crew shifts of the train-drivers across the entire country. This is a massive scheduling problem; and our instances feature graphs with billions of edges. We could solve these problems with high quality, i.e., optimality gap usually around 1-2%, consistently and within few days of computation time. Our largest instance has 2.3 billion edges and was solved to an accuracy of 4.74% within 104 hours. This was a substantial breakthrough, since the methods applied previously did not scale accordingly.

Our algorithm solves a combinatorial optimization problem called SET CONSTRAINED FLOW, by employing the column generation method in a modular software architecture. SET CONSTRAINED FLOW generalizes MINIMUM COST FLOW: It allows to specify arbitrary sets of edges, and for each such set two numbers, a lower bound and an upper bound. Any feasible solution must provide total flow in each such set between the required bounds. Our main design goals and decisions have been: We require that the solver is applicable to combinatorial optimization problems, with reasonable modeling efforts in terms of SET CONSTRAINED FLOW. Many real-world use-cases have *locality* and *globality* properties, that shall be supported in the modeling: There is a clear separation between groups of constraints that link resources at a narrow perspective, e.g., precedence requirements for tasks, and groups of constraints that have to hold at a high level view, e.g., task coverage requirements. The solver has to scale "out of the box" with column generation, which iterates between solving master- and pricing-problems, and where we solve the pricing-problems in parallel. The software can be accessed easily through a REST API, which hides the technical deployment details from a user.

The main module of the solver is the *coordinator*: It instructs the behavior of the remaining modules, observes if a terminating solution has been found or if time-limits have been exceeded. Following the column generation approach, the constraints of the instance are split into local- and global constraints, respectively. The *master solver* hosts the global constraints and always solves some given linear program. The corresponding dual solution yields reduced cost coefficients on the edges, which are reported to the pricing solver by the coordinator. We currently have connectors to CPLEX, Gurobi, HiGHS, and SCIP to implement the solution of the linear programs. The *pricing solver* solves the subproblem, which is given by the graph, the local singleton constraints, and the vertex-balances. In many cases, these problems are simply SHORTEST PATH problems of some variety (and determined by the instance classifier). The pricing solver is responsible for most of the speed and scalability of our approach: Especially on directed acyclic graphs, SHORTEST PATH can be solved in linear time and a large number of paths can be generated in parallel. The paths are reported to the *column pool*, deciding which ones to keep and which

ones to discard. The remaining columns are finally added to the linear program to be solved by the master solver. Finally, the *fixing solver* decides which variables or constraints in the linear program are rounded, frozen, or released. This module can implement branch-and-bound schemes (if need be) or fixing heuristics originating from insights into the use-case at hand.

Crew Scheduling

Our optimizer is used for annual driver- and conductor-planning and for strategic simulations at the Swiss Federal Railways. The CREW SCHEDULING problem asks to partition a set of tasks into a set of shifts, while respecting temporal consistency, geographic consistency, break requirements, and other labor rules. Drivers are organized into depots, where each depot has one or more driver groups with different skill sets. The objective is to minimize a weighted sum over the shifts. We applied the SET CONSTRAINED FLOW approach as follows: The graph captures the possible transitions between the tasks. States stored on the vertices represent a task and the relevant history of a shift that contains it. A source vertex marks the start of a planning day at some location, while a target vertex marks its end there. By construction, each path from the source to the target represents a valid shift. Set constraints ensure that at most one unit of flow passes through each vertex. Hence any integral flow decomposes into disjoint paths.

Our reference instances cover a variety of real-world scenarios, ranging from single-depot to full-country optimizations. The largest instance, covering the entire Swiss-wide optimization problem, has 20041 tasks, 67 overall driver groups, yielding a graph with 2.3 billion edges, which was solved within 104 hours of compute time to an accuracy (integrality gap) of 4.74%. The overall results substantiate that our technology is capable to solve very large instances in high quality, consistently.

Biography

PD Dr. Alexander Souza is the founder and CEO of Algomia. He has many years of professional experience in the optimization and software industry. He holds a doctorate degree from the ETH Zurich and holds a lectureship on algorithms at the University of Zurich.

Telco AI Assistant: From research to delivery – challenges and lessons learned

Iwona Wojdan, Marcin Ziółkowski

Orange Innovation Poland

Iwona.wojdan2@orange.com, Marcin.ziolkowski@orange.com

TAIA (Telco AI Assistant) is an AI-driven agentic system that uses advanced reasoning to automate the troubleshooting process for L3 connectivity issues. TAIA's specialized agents work together to diagnose the issue and provide actionable insights to help administrators remediate the cause.

In this talk, we demonstrate how agentic systems can address real-world telecommunications challenges, with a focus on root cause analysis of L3 connectivity issues, customer support automation, and operational efficiency optimization. We will present the L3 connectivity troubleshooting workflow as a concrete Operations Research use case, showing how AI-based reasoning in a multi-agent system is applied to a real production problem, delivering value through significant reduction in the time required to resolve non-functional network paths and consequently lowering the costs associated with deploying and migrating new systems within the network. TAIA is deployed as a module within IPManager, a production network-management platform used at Orange Polska, which supports automated operations across multiple network domains with improved resilience and resource utilization.

TAIA extends IPManager with AI-based reasoning that automatically gathers data, identifies root causes, generates clear problem descriptions, and provides actionable recommendations, replacing manual, user-driven troubleshooting workflows. We will share how this use case evolved from a research prototype into an operational product and how to increase Technology Readiness Level of the research proof of concept. We will illustrate key challenges: how to build trust and gain visibility for a research solution in an operational environment by establishing effective and efficient collaboration between multiple teams with different areas of expertise.

Finally, we will share how the collaboration between research and operational teams has benefitted both: with faster prototyping and knowledge transfer for the operational team, and increased quality assurance revealing tricky edge cases in operational processes which the research Proof of concept did not envision, highlighting what had to change for a research-driven solution to succeed in a real operational environment.

Biography

Iwona Wojdan: I work at Orange Innovation Poland in the Virtual and Programmable Infrastructure competence center, where I'm an Expert R&D engineer. My focus is on building and evolving programmable, cloud-native network platforms – from virtualization layers all the way up to how services are orchestrated and exposed. A big part of my work is making infrastructure behave more like software: automated, testable, and scalable.

Before that, I was part of CodiLime's Professional Services team, working as a Network Engineer in areas like SDN, NFV, Linux, and Open-Stack. I was directly involved in designing and delivering software-defined network solutions for operators and large customers, often bridging the gap between low-level networking and modern cloud infrastructure.

Earlier, I worked at Orange Labs Poland in the International and Interdomain Projects division. There I was an expert working across multiple countries and domains, helping to align architectures, platforms, and operational models between different Orange organizations. That experience taught me how hard distributed systems really are – not just technically, but also organizationally.

Before moving into cloud and virtualization, I spent many years at Orange Polska in the Aggregation Networks Development division as a Main Network and Service Architect. I was responsible for designing large-scale carrier networks and the services running on top of them. That background in real-world telecom networks strongly shapes how I think about reliability, scale, and failure today.

All of this together is what I bring to Kubernetes and cloud-native infrastructure: a mix of deep networking, large-scale production systems, and modern programmable platforms.

Marcin Ziółkowski: As an R&D Expert at Orange Innovation Poland who is also pursuing a PhD at WUT, my background bridges the gap between academic research and industrial applications. With years of Software Engineering and research experience, I focus on scalable, efficient, and sustainable distributed systems, pushing projects from conceptual research to production.

Holding all 5 Kubestronaut certifications (CKA, CKAD, CKS, KCNA, KCSA) and years of experience running Kubernetes clusters under my belt, I apply insights from the Hexa-X-II (EU 6G Project) to my ongoing research in cloud-native Telco and Open-Source technical enablers.

Unlocking network-wide profitability steering – the case of automated network initialization

Sander Van Aken

Flix SE, Warschauer Platz 11-13, 10245 Berlin, Germany
sander.van-aken@flix.com

Flix is global player in the long-distance passenger transportation market, with networks in 46 different countries across 5 continents with brands as FlixBus, FlixTrain, Greyhound (US) and Kamil Koc (Turkey). As a commercial company, tailoring supply to demand is a key element for sustainable growth and profitability for both Flix and its 1000+ partners operating the network. Our network planners create timetables, bus and driver schedules with the aim to build these profitable networks. Planners anticipate and adapt to highly seasonal demand in a strongly data-driven manner.

Our tech teams have developed a range of ML- and OR-based tools to enable building a network of routes to maximize profitability. Due to the vast and interconnected nature of the network, this requires accounting for complex cross-line demand interactions, often referred to as the "network effect". Properly forecasting profitability requires considering the complete network. However, the yearly planning process happens more than a year ahead of operation. Till now, this meant to an ever-increasing network maturity over time and not all information being available at the early stages. This poses major challenges to effectively utilizing these tools early on and at scale.

To address this, we developed our Automated Network Initialization solution, combining ML-based demand seasonality forecasting, OR-based supply optimization, and business domain knowledge. By optimizing the supply for the anticipated year-on-year demand changes, we created a complete seasonalized network at the start of the network planning process. Developing this was far from a "textbook" optimization task. Navigating the "chicken-and-egg" problem of demand and supply interaction, developing an optimization model without a proper quantification of what "good" means, absence of benchmarks, and dealing with many special cases and data sources. This required a multidisciplinary approach in a strong cross-team and business-tech collaboration to deliver impactful and trusted results.

The results have been transformative. We successfully initialized 1,400 lines for the 2027 network automatically, covering approximately 93% of our mature markets. Consequently, planners can now focus their attention on lines where improvements yield the highest return on invest, enabling network-wide profitability steering from day one. In this talk, we will share our journey developing this solution and how we successfully addressed the many challenges encountered along the way.

Biography

Sander Van Aken is Staff Operations Research Engineer at Flix SE (known from FlixBus, FlixTrain, Greyhound US and Kamil Koc). As part of the network planning optimization team, he designs, develops and tests custom-built OR-algorithms for timetable, bus and driver scheduling, in close collaboration with both tech- and business stakeholders. These get embedded in a software product owned and developed by the team. Sander holds a MSc. in Traffic and Logistics Engineering and prior to joining Flix worked as advanced analytics consultant mainly in the public transportation industry on simulation-, data science and optimization-based projects.

Computational modelling of a carrier assignment problem under the ROAR-NET API

Carlos M. Fonseca, Ulisses M. Rodrigues

University of Coimbra, CISUC/LASI, DEI, Portugal

cmfonsec@dei.uc.pt, ulissesrodrigues@student.dei.uc.pt

The NEXUS Agenda consortium is led by the Port of Sines and comprises 35 partners, including port authorities, maritime and terminal operators, railway operators, carriers, dry ports, logistics operators, technology suppliers, importers, exporters, universities and research institutes. Within the NEXUS framework, Shipperform is a service procurement platform under development by Devlop that will allow shippers to request and obtain tailored quotes for their shipping needs directly from selected carriers and to determine and contract the winning carriers, among other features.

In the tendering process of interest, shippers define lanes in terms of origin, destination, and load volume. Carriers bid for the lanes of interest to them by quoting a price for the load volume that they are willing to transport in each lane. Carrier assignment consists in determining one or more winning bids for each lane so as to simultaneously minimise total cost and unassigned load volume, which is a multiobjective optimisation problem, subject to shipper-specific constraints such as minimum and maximum numbers of carriers selected and minimum and maximum load per carrier. The shipper may then mark some of the lane-carrier assignments as final and start a new tendering round to fulfil the remaining needs.

Rather than developing a dedicated carrier-assignment algorithm, a local-search model of the problem was developed and implemented in a solver-independent way based on the ROAR-NET API [1], which allows solvers implemented using the same API to be applied to the model directly in an off-the-shelf fashion. In particular, a preference-aware version of Pareto Local Search (PLS) can be used to explore the price versus unassigned load trade-off with the goal of presenting alternative assignments to the shipper to support their decision making. The presentation will focus on the development of the problem model and preference-aware PLS solver under the ROAR-NET API with a view to their integration in Shipperform. Illustrative examples and experimental results based on artificial instances will also be presented.

Acknowledgement This work was supported by Portuguese Recovery and Resilience Plan (PRR) within the scope of the Agenda "NEXUS: Pacto de Inovação-Transição Verde e Digital para Transportes, Logística e Mobilidade" (Investment Project No. 53) under Grant C645112083-00000059.

Reference

[1] The ROAR-NET API Specification. <https://github.com/roar-net/roar-net-api-spec>

Biography

Carlos M. Fonseca is an Associate Professor in the Department of Informatics Engineering of the University of Coimbra, Portugal, and a member of the Adaptive Computation Group of CISUC, the Centre for Informatics and Systems of the University of Coimbra. He obtained his doctoral degree from the University of Sheffield, U.K., where he also conducted postdoctoral research. His main research interests are in multiobjective optimisation, evolutionary computation, experimental evaluation of optimisation algorithms, and practical applications of optimisation. Since 2023, he has served as Action Chair and Grant Holder Scientific Representative of COST Action Randomised Optimisation Algorithms Research Network (ROAR-NET), a European research and innovation network aimed at making randomised optimisation algorithms widely competitive in practice by identifying and reducing obstacles to their adoption at the scientific, technical, economic, and human levels.

StudyPlanner: Helping students to plan university courses with integer linear programming

Victor Bosneag, Utku Dündar, Kai Eberl, Zixuan Fan, Ahmed Mhadhbi, Michael Ritter, Justus Schneider, Anthony Tang, Philipp Wiedmann, and Andreas Wiese

Technical University of Munich, School of Computation, Information and Technology, Department of Mathematics
Boltzmannstraße 3, 85748 Garching bei München, Germany
andreas.wiese@tum.de

At the beginning of each semester, all university students need to plan the courses they want to take in this semester. In the process, they need to take into account all rules of their degree program. Those might, for example, regulate that some courses are compulsory, that from some groups of courses one needs to select a certain number, or that one needs to select a study focus from some given options. Also, there are additional considerations, such as the preferences of the student for certain subjects, the times when the courses are scheduled, and that some courses might need to be taken in a certain order. In particular, it is a good idea to not just plan the courses of the upcoming semester, but to plan all courses in all future semesters.

We developed the StudyPlanner, which is an optimization tool that students at the Technical University of Munich (TUM) can use in order to plan their courses, taking all the considerations above into account. Its backbone is a mixed-integer programming (MIP) formulation that balances multiple objectives, including, for example, to minimize the time until graduation and to maximize the satisfaction of the student with the selected courses. We can solve it within seconds using the open source MIP-solver HiGHS. In this talk, we will discuss technical aspects of modeling the optimization problem suitably and also highlight important points that arose in discussions with relevant stakeholders at TUM. As of now, the StudyPlanner supports 15 study programs at TUM, and it is publicly available at <https://studyplanner.co.cit.tum.de>.

Biography

Andreas Wiese studied mathematics at the TU Berlin where he also did his PhD in mathematics. He was a postdoc at the "La Sapienza" University in Rome and at the MPI for Informatics in Saarbrücken. Later, he became a professor at the Universidad de Chile in Santiago and at the Vrije Universiteit Amsterdam. Finally, Andreas Wiese joined the Technical University of Munich (TUM) in 2022. In his research, he works on discrete optimization problems in theory and practice.

Developing an optimization tool for farm advisory

Rasmus Bang¹, Julio Goez^{1,2}, Mario Guajardo^{1,2}

¹SNF – Centre for Applied Research at NHH, Norway

²NHH Norwegian School of Economics, Norway

mario.guajardo@nhh.no

This talk will present an overview on an on-going project in Norway called OPTINORFARM. The project proposes to implement optimization models and methods into computational tools to support decision-making of farmers. We develop a quantitative framework to economically optimize production at farm level under various operational constraints and to calculate the corresponding greenhouse gas (GHG) emissions. The economic optimization is carried out by a mathematical programming model. The assessment on GHG emissions is carried out using well-established methods to compute methane, nitrous oxide and carbon dioxide (CO₂) emissions in CO₂ equivalent.

Our framework can be used to study a variety of different problems in cattle farming. For example, we may address questions such as: Is it optimal for the farmer to operate their farms under conventional or organic systems? Which measures should farmers take in order to be economically profitable and climate friendly?

Such types of questions are often faced by farm advisors, who work closely with farmers to assist them in how to run their farms. In particular, the debate between economic profitability and environmental impact is currently a hot debate in policy making. Although policies shape goals and provide guidelines, the profitability of dairy and beef products and the environmental impact of their production depend on decision-making at cattle farm level. Thus, it is important to identify farm-level actions which contribute to the economic profitability of farmers, while at the same time mitigate the impact on the environment.

Our optimization tool is developed in close cooperation with farm advisors. The challenges have not only been about technicalities in regards to model formulation, but also about how to implement the model, how to communicate its capabilities, and how to create competences among practitioners.

Biography

Mario Guajardo is Professor in Operations Research at NHH Norwegian School of Economics. He is currently project manager of OPTINOR-FARM, a collaborative project between academia and industry about the economic and environmental optimization of Norwegian farms.

Julio Goez is Associate Professor at NHH and participates in OPTINORFARM as researcher and work package leader.

"Believe me, it is optimal": Designing online optimisation systems for explainability & acceptance – A case study

Christopher Bacher

Zalando SE, Germany

`christopher.bacher@zalando.de`

In this talk we have a look at how Zalando developed its current Online Order Allocation algorithm which makes hundreds of millions of live routing choices for customer orders every year and which design decisions have been made to ensure stakeholder acceptance and support explainability.

When changing or replacing an(y) existing system, it must not only meet purely functional requirements but also non-functional ones to ensure that it is in the end accepted & adopted by stakeholders & business decision makers. This is true more so if such a system contains black-box components whose details are not easily inspectable or understandable. Optimisation algorithms – even when not using black box solvers – often fall in this category for business stakeholders.

We detail how the algorithmic design, such as choice of objective function and preprocessing steps, can contribute to explainability and drive stakeholder acceptance. Further we discuss which support systems & metrics can ease inspection & debugging of a live online optimisation algorithm as well as help to answer common stakeholder questions. Finally we outline how we approach experimentation, deal with difficulties in A/B testing, and how we made our system fit for discrete event simulation to allow powerful what-if analyses.

Biography

Christopher Bacher works as a Principal Applied Scientists at Zalando SE. In the past years he has contributed to redeveloping Zalando's order allocation systems, improving Zalando's delivery time predictions, and improving Zalando's experimentation methods in the fulfillment domain. Before he worked on the relationship between formal grammars & dynamic programming, as well as on the optimization of electro-hybrid powertrains.

Building an optimisation platform for everyone

Justus Kilian Wolff

Kleiner Schäferkamp 40, 20357 Hamburg, Germany

`justus@bools.io`

Small and medium-sized businesses face complex scheduling, allocation, and logistics challenges daily – problems that mathematical optimisation solves elegantly. Yet most SMBs do not use these tools. In fact, most do not even know they exist. "Operations Research" as a term means little to the average business owner; they often have no mental model for what it offers. We believe this is partly our own doing as a field: we have made optimisation too complex, too specialised, too hidden behind solver syntax and mathematical notation. However, the rise of large language models has changed what is possible. Topics that previously required deep technical expertise are now becoming accessible to broader audiences through conversational interfaces – and we believe optimisation is ready for this shift.

We set out to build bools, an optimisation platform designed to make OR accessible to everyone. At its core is a declarative modelling approach where models read like business requirements rather than solver code. This readability is the key differentiator: users can understand, verify, and modify their models without OR expertise. The platform caters to multiple levels – OR experts can use the bools library directly for full control, while employees without technical background can formulate problems through conversation with an LLM-powered assistant. Both paths produce the same readable model, translating complex optimisation into clear, business-friendly terms.

In this talk, we present the end-to-end workflow of the platform: from a user describing their problem in natural language, through the LLM assistant translating that description into a structured optimisation model, to solving and presenting results back in terms the user understands. We demonstrate how the platform's architecture work together to lower the barrier from problem recognition to deployed

solution. Along the way, we share challenges and lessons learned, particularly around integrating LLMs as an assistant layer: where they work well, where they struggle, and what guardrails are needed to produce reliable models.

We conclude with reflections on collaboration and language: how do OR specialists work effectively with end users who have never heard of the field? Our experience suggests that framing matters – "decision intelligence" or "better scheduling" often resonates more immediately than "Operations Research". We believe that as better tools lower the barrier to entry, more people will experience the value of optimisation firsthand, and the field will naturally gain the recognition it deserves.

Biography

Justus Kilian Wolff is the founder of bools, building a platform to make mathematical optimisation accessible to everyone. He holds a BSc in Industrial Engineering from Karlsruhe Institute of Technology and an MSc in Operational Research from the University of Edinburgh. After building fraud detection systems, he decided to go out on his own and build an optimisation platform that bridges the gap between complex models and the people who need them.

Delivering optimization to non-experts with structured dialogue, dataset profiling, and solver verification

Matthew Jurado, Peter Wolf

Independent Researcher, Germany

juradomatt@yahoo.com, pe.wol@mailbox.org

Operations Research projects often fail to reach end users because requirements arrive as partial narratives, assumptions remain implicit, and outputs are hard to reproduce or explain in business terms; recent work on learning for combinatorial optimization and explainable AI for OR reinforces the need for transparent, auditable delivery pipelines (Bengio et al., 2021; De Bock et al., 2024; IDEAS/RePEc). We present an end-to-end delivery workflow that converts natural language stakeholder questions into verified optimization runs and stakeholder-ready explanations. A large language model (LLM) is used only for dialogue and orchestration: it clarifies intent, asks targeted follow up questions, and produces a validated decision specification. Deterministic tools then bind the specification to a canonical data model, a standardized schema of tables and keys, and execute template optimization models; dataset onboarding computes a short, cached profile of schema coverage and value domains keyed by a dataset fingerprint to ground clarifying questions and defaults. Reproducibility is maintained via versioned model templates, recorded decision specifications, dataset fingerprints and profiles, and solver-side verification, such as feasibility checks and objective recomputation. We share early implementation lessons, including failure modes such as ambiguous intents, missing parameters, brittle structured outputs from the LLM, and user confusion when exposed to technical settings, along with mitigations including bounded clarification, strict validation, robust structured output parsing, and progressive disclosure in the interface. Results surface explicit assumptions, key cost drivers, and binding constraints, and we conclude with reusable artifacts for OR teams: decision spec templates for stakeholder intake, a dataset onboarding and profiling workflow, and a verification and explanation checklist for trustworthy end-to-end delivery.

References

Bengio, Y., Lodi, A., & Prouvost, A. (2021). *Machine learning for combinatorial optimization: A methodological tour d'horizon*. European Journal of Operational Research, 290(2), 405–421.

De Bock, K. W., et al. (2024). *Explainable AI for Operational Research: A defining framework, methods, applications, and a research agenda*. European Journal of Operational Research, 317(2), 249–272.

Biography

Matthew Jurado is an operations research practitioner and independent researcher focused on optimization and simulation for supply chain and fulfillment operations. He is currently the Head of Operations Research for a fast-growing grocery e-commerce company in Central Europe. He began his career in finance and later transitioned into warehouse science, specializing in discrete event simulation and mathematical optimization. He built and led an operations research team at a Fortune 500 company in the USA. After relocating to Europe, he led a tactical planning team that applied optimization techniques to short-, mid-, and long-term planning horizons. He is represented here as an independent researcher.

Peter Wolf is an operations research engineer currently working at a grocery e-commerce company, with a focus on operations planning and strategy. Previously, he worked as an operations research specialist at one of Europe's leading pet supply e-commerce companies. He began his career as a research assistant at the Chair of Logistics and Quantitative Methods at the University of Würzburg in Germany. During his time there, he coauthored research at the intersection of machine learning and optimization and developed machine learning and optimization models for a local temperature-controlled transport provider and capacity planning models for an aviation spare parts and maintenance services company. He is represented here as an independent researcher.

Bridging the gap in academia-industry collaboration: Research competitions as a catalyst for open innovation

Matteo Pozzi

Optit srl, Via Mazzini 82, 40138 Bologna, Italy
Matteo.pozzi@optit.net

In the current landscape, internal R&D teams often struggle to keep pace with the rapid evolution of complex business challenges. Consequently, large corporations are increasingly adopting **Open Innovation** paradigms to tap into external expertise from researchers, startups, and scale-ups.

A highly effective instrument in this domain is the **Research Competition**. By providing a detailed use case and realistic datasets, corporations can invite external parties to propose innovative solutions. Beyond the immediate monetary award, these competitions serve as a strategic entry point for long-term collaborative partnerships.

This approach was a cornerstone of the **Horizon Europe project TUPLES (Trustworthy AI)**, which featured a challenge proposed by **Airbus**. A key finding from this experience was the critical lack of affordable, specialized platforms capable of evaluating **how models work** (their internal logic, robustness, and trustworthiness) rather than just comparing final numerical results—a limitation found in the most popular players, like Kaggle.

To fill this gap, a novel **Competition Platform** has been developed, to balance technical rigor with strategic management:

Technical Excellence & Reproducibility:

- Full Environment Parity via Docker: The platform orchestrates isolated, high-performance containers, ensuring that submitted models run exactly as they did in the participant's local environment, eliminating the "it works on my machine" problem.
- Beyond Black-Box Leaderboards: It moves beyond numerical scores by providing tools to analyze constraint satisfaction, model robustness, and decision logic, ensuring solutions are not just accurate, but operationally viable.
- Standardized Benchmarking: It provides a robust framework to compare diverse approaches (from MILP to Metaheuristics) under identical computational resource constraints (CPU/RAM).

Strategic Management & Ecosystem Building:

- White-Label Business Model: Organizations can launch proprietary, fully branded challenges, maintaining complete control over the user experience and their intellectual property.
- Private Ecosystems: The platform allows for the use of proprietary datasets in a secure environment, enabling corporations to cultivate a dedicated community of research partners and tech providers.
- Trustworthy AI & Explainability: By integrating advanced evaluation metrics and GenAI/LLM-based explainability tools, the platform helps bridge the gap between academic research and industrial adoption.

The Platform is now fully operational. By presenting its genesis and development, we aim to introduce the OR practitioners' community to a powerful, untapped tool for driving effective, data-driven Academy-Industry collaboration.

Biography

Matteo Pozzi, following a MSci in Physics and a Diploma in International Relations, started a career as Management Consultant in Italy and the UK.

Partner and CEO of Optit since 2010, he has been leading the company (spin-off of the Operations Research team of the University of Bologna) from start-up phase to becoming a leading Italian player in the development of Decision Support Systems based on forecasting, data analytics, simulation and optimization models, with significant operations across the EU and the USA.

Internal transport optimization opportunities for warehouses

Paweł D. Domański¹, Szymon Cyperski², Artur Sobas², Piotr Maciejewski²

¹Warsaw University of Technology, Institute of Control and Computational Engineering,
Nowowiejska 15/19, 00-665 Warsaw, Poland
pawel.domanski@pw.edu.pl

²Betacom SA, ul. Połczyńska 31a, 01-377 Warsaw, Poland
scyperski, asobas, pmaciejewski@betacom.com.pl

Simulation constitutes a backbone during design and operation of logistics warehouses. Agent-based modeling (ABM) framework enables to introduce human operation specifics into the system. Once the model uses human operation, than the optimization may also take it into consideration. Distributional warehouses operate in complex non-stationary supply chain system. Their operation is crucial for other elements and should not limit their performance. Their operation should be transparent, predictable, repeatable and at least close to optimal. Warehouse operation uses many process such as orders routing, batching and scheduling, picker assignment, batches networking, and location assignment. This work addresses the assessment of such optimization tasks for full scale warehouse. The results are compared with the current practice generated by classical Warehouse Management Systems (WMS). It is shown that custom optimization allows to improve warehouse operation and brings quantitative benefits.

Biography

Paweł D. Domański was born in 1967 in Warsaw, Poland. He received the M.Sc. degree in 1991, the Ph.D. degree in 1996 and the D.Sc. degree in 2018, all in control engineering, from Warsaw University of Technology. Since 1991 he has been working at the Institute of Control and Computational Engineering, Warsaw University of Technology. Besides scientific research, he has participated in dozens of industrial implementations of Advanced Process Control (APC) and optimization. He is the (co-)author of four books and more than 160 publications. His main research interest is in industrial control and optimization, and multi-agent modeling for the supply chain management.

Digital twins for simulation and optimization of heavy equipment multipart logistics

Rohit Gupta, Aniket Mitra

Computationele Wetenschapsgroep

rohit@computationelewetenschapsgroep.com, aniket@computationelewetenschapsgroep.com

Transport of heavy equipment (for example: parts of a wind mill, fabricated sections of a house, parts of an aircraft) is primarily constrained by factors such as (a) the cost of (multi-modal) transportation / storage of the equipment (b) a predetermined assembly sequence. This gives rise to a logistical puzzle of Just-In-Time delivery where the right equipment part has to arrive at the right time & in the right sequence while minimising transport, inventory & project overrun costs. Replicating planning & optimisation digitally can enable planners to simulate scenarios & find optimal solutions. To achieve this we need to model the physical process digitally by (a) creating a unified semantic model of assets, schedules & constraints (b) enabling planners to run simulations and perform stage-wise optimisation. In this talk we will present a system that allows users to run simulations using a Discrete Event Model (DEM) on a digital twin. The system further allows for simultaneous simulation & optimisation by coupling DEM with optimisation techniques like simulated annealing. Furthermore, we demonstrate how such a logistics workflow can be modelled with a Planning Domain Definition Language (PDDL) that enables domain experts to easily configure / modify the workflow in-order to run simulations & identify optimal solutions.

As a subsequent extension of this effort we intend to further explore the automation of planning problems using PDDL primitives and variants of reinforcement learning. Using PDDL for describing the process formalizes it and provides it a logically connected structure. This step is important when we have a digital twin that is powered by operational data. The presence of a structure (in the form of PDDL) grounds the information collected, as the process evolves and planning is executed. This grounding can provide the basis for a generative AI based conversational interface that can enable finer control and reconfiguration of the process in situ. The concept is powerful for replanning because of disruptions during operations and further for designing around business cases that are defined in natural language and have tangible KPIs that businesses steer upon.

Modelling planning problems as graphs with dependencies and relationships can lend itself to high reward sequences which when accurately encoded significantly accelerate learning of a policy that follows problem constraints and adapts to unseen distributions (of states

and actions). Reinforcement learning benefits in terms of reward shaping and finding a balance between exploration and exploitation when specific sequence of actions can be traced to desirable end states.

Finally combining the strengths of OR techniques with recent advances in AI grounded by a symbolic system has the potential for widespread adoption and value creation.

Biography

Dr. Rohit Gupta is a Director at Computacionele Wetenschapsgroep (CW). Rohit has been at the forefront of research and engineering for the past 20 years. For the past decade, Rohit has been using his background in computational sciences to build computational models for a number of companies including Nvidia, ASML, KLM & HERE. Rohit holds a PhD from TU Delft, has published in numerous scientific journals and holds multiple patents.

Aniket Mitra is the Founder of Computacionele Wetenschapsgroep (CW). Aniket's expertise lies in building semantic models that provide domain specific symbolic representations for simulation & optimization models.

Panel discussion

Overcoming resistance to change – integrating OR models into existing business processes

Unlocking the impact of OR often requires integration into – or changing – existing business processes. With our panelists, we will explore challenges, lessons learned, and practical strategies for navigating operational friction and overcoming human resistance.

Moderation: Sander Van Aken (Flix SE)

Panelists:

- **Laura Houlley** is R&D Project Manager at Air Liquide, where she leads digital projects optimizing the bulk liquid supply chain. Her role involves not only managing internal R&D efforts but, critically, coordinating interfaces with operations and the internal entities responsible for industrializing these tools. Her background includes designing schedule optimizers and machine learning forecasts, as well as leading product development in supply chain solutions.
- **Victoria Guerrero** is a Senior Data Scientist at Ingka Group (IKEA), specializing in operations research and data-driven decision making. She began her career working on the integration of renewable energy into electric power systems and, over the past six years, she has focused on improving logistics processes through analytics and optimization. Beyond her technical work, Victoria is passionate about the intersection of business, product, and decision science. She enjoys engaging with different stakeholders to increase adoption and make onboarding easier, while translating real needs into meaningful product improvements.

Discussion group topics

OR + GenAI – how they complement each other

Room 4.01

Sofiane Oussedik (IBM), Carl Johan Kellerth Fredlund (Jeppesen ForeFlight)

Are we moving toward a world where a business manager with a prompt-engineered agent replaces a specialized OR practitioner?

This discussion group is dedicated to everyone interested in learning about generative AI and leveraging its capabilities to support OR projects. We'll discuss and share experiences on how GenAI can be leveraged in OR projects to speed-up the development of OR applications and how it can support business users. We'll also exchange about the associated benefits and risks.

Practical struggles of getting solutions into production

Room 4.02

Björn Thalén (B3 Indes) , Torkel Haufmann (Sintef),

When the end goal is to create a difference, in practice it is not always enough to create a good solution. Sometimes the problem lies in technical implementation of a scaled up solution, other times it is organizational where the proposed end users are not as keen as whomever initiated the project. In this discussion group we share experiences around the struggles we have had and if we found a solution to it, or not.

Academia vs. Industry: The relevancy gap

Room 4.04

Vladimir Fux (Zalando), Sander Van Aken (Flix SE)

Novelty vs Impact:

- Academia rewards "Novelty." You cannot get tenure by showing that a 40-year-old algorithm works perfectly on a new dataset. This forces researchers to invent "Frankenstein models"—unnecessarily complex hybrids—just to get published.
- Industry rewards "Impact", which is largely biased towards existing, simple, well performing tools and methods that are guaranteed to work. A practitioner would rather use a sub-optimal model they can explain to a VP in 30 seconds than a state-of-the-art decomposition method that breaks if the data is 5% noisy.

Theory vs Real World:

- Academic work often focuses on achieving optimality and other theoretical properties in sterile environments. Many OR graduates are highly trained in theory and solution-first thinking. As a result, they might struggle with framing messy business problems, working with users, clients, and delivering solutions in agile, iterative, and MVP-styled approaches. Which knowledge/approaches are the most relevant then for a fresh OR graduate, transitioning to the industry?

Maps

Conference venue:

The conference will take place at the Warsaw University of Technology (Politechnika Warszawska), in the Rektorska4 building.

Address: Rektorska 4, Warsaw University of Technology, 00-664 Warsaw, Poland

Google Maps link: <https://maps.app.goo.gl/nYaH66DDZb8QoDbW9>



Conference dinner:

Restaurant DOCK19, Wybrzeże Kościuszkowskie 43C.

The restaurant is located between the Copernicus Science Center and the Elektrownia Powisle shopping center.

GPS: 52°14'28.4"N 21°01'41.6"E

Google maps link: [dock19](#)

From the conference venue, you can take the metro line M1 from Politechnika station (direction Młociny) and go to Światokrzyska station (two stops), then change to the metro line M2 (direction Brodno) and get off at Centrum Nauki Kopernik (Copernicus Science Center) station (two stops), and then walk 270m.

From the MDM hotel, you can walk to Politechnika metro station (see above). Alternatively, you can take the tram in front of the hotel heading to Światokrzyska, get off at Metro Światokrzyska, and then take the metro line M2 (see above).

A 20-minute ticket should be sufficient.