

Dutch Day on Optimization 2026

Groningen

PROGRAMME BOOKLET

Programme at a glance

10:00–10:50	Welcome with coffee and tea
10:50–11:00	Introduction
11:00–11:45	Plenary presentation — Bernardo Pagnoncelli
11:45–12:30	Contributed talks
	Steven Miltenburg (<i>VU</i>) — <i>On the complexity of capacitated vehicle routing with order restrictions</i>
	Koen Ligthart (<i>TU/e</i>) — <i>Periodic convexity of integer program value functions and applications to block-structured integer programs</i>
	Yingzhu Liu (<i>RUG</i>) — <i>Distributionally Robust Federated Learning with Multi-Source Data</i>
12:30–13:45	Catered lunch
13:45–14:30	Plenary presentation — François Glineur
14:30–15:15	Contributed talks
	Wenhao Chi (<i>UvA</i>) — <i>Violation Certification for Linear Predict-Then-Optimize Pipelines</i>
	Mathijs Barkel (<i>TiU</i>) — <i>Integer semidefinite programming for the maximum k-colorable subgraph problem</i>
	Gijs Bakker (<i>RUG</i>) — <i>A dynamic overbooking model for multi-compartment parcel locker systems</i>
15:15–15:45	Coffee break
15:45–16:30	Contributed talks
	Lucas Ketels (<i>RUG</i>) — <i>Convergence Rates for Stochastic Proximal Methods via Regularized Optimality Gaps</i>
	Daan Noordenbos (<i>TiU</i>) — <i>Personalised versus Posted Pricing from Samples</i>
	Wick Wijnholds (<i>UT</i>) — <i>Optimal Antenna Design: An Exact Approach</i>
16:30–17:15	Plenary presentation — Nicole Megow
17:15–18:00	Closing & drinks
18:00–21:00	Dinner

Keynote speakers

Bernardo Pagnoncelli

SKEMA Business School

BRIDGING MARKOV DECISION PROCESSES AND MULTI-STAGE STOCHASTIC PROGRAMMING

Abstract. Multi-stage stochastic programming (MSP) and Markov decision processes (MDPs) provide two complementary frameworks for sequential decision-making under uncertainty. In this talk, I will discuss a modeling framework that brings key features of MDPs into MSP. Using policy graphs, the framework distinguishes between continuous physical states, describing the evolving system, and discrete regime states, describing the stochastic environment. This representation accommodates both finite- and infinite-horizon problems while retaining the optimization structure of MSP.

I will then discuss two extensions. First, transition probabilities between regimes may depend on decisions, allowing actions to affect the future stochastic environment. Second, partially observable models allow beliefs about the underlying system to evolve through Bayesian learning. Together, these extensions capture settings in which decisions can influence both future uncertainty and the information available about it.

The resulting models can be solved approximately using extensions of stochastic dual dynamic programming (SDDP), providing an alternative to traditional discretization-based approaches for structured MDPs with continuous state and action spaces. I will illustrate the framework through examples and applications, including ongoing work on sequential decision-making in forestry under climate change.

Short bio. Bernardo Pagnoncelli is a Full Professor at SKEMA Business School in France. His research focuses on stochastic optimization, particularly stochastic programming and sequential decision-making under uncertainty. He develops models, algorithms, and computational methods for optimization under uncertainty, with a particular interest in multistage problems and risk-averse decision-making. His work is motivated by applications in several areas, including finance, mining, forestry, energy, and transportation. His research has been published in leading journals in operations research and optimization, including *Operations Research*, *Management Science*, *Mathematical Programming*, and the *SIAM Journal on Optimization*. He holds a PhD in Applied Mathematics from PUC-Rio.

François Glineur

UCLouvain

PERFORMANCE ESTIMATION OF OPTIMIZATION METHODS: AN OVERVIEW

Abstract. Identifying the worst-case behaviour of an optimization method can itself be cast as an optimization problem. This is the main idea behind performance estimation, initially proposed by Drori and Teboulle in 2014. In this framework, one seeks to compute the exact worst-case performance of a given black-box optimization algorithm after a prescribed number of iterations over a given class of problem instances. In many cases, including most first-order algorithms with prescribed coefficients, this computation can be reformulated as a finite-dimensional, tractable semidefinite program using necessary and sufficient interpolation conditions for the function class involved. Primal and dual optimal solutions to this program provide a tight worst-case bound, a mathematical proof establishing this bound and a problem instance where this worst-case behaviour is achieved.

Knowledge of the exact worst-case behaviour of algorithms can be used to compare efficiency across methods, tune algorithmic parameters (such as step sizes) and, ultimately, design new methods with improved guarantees, either using insights gained from the performance estimation procedure or using automated method design techniques. In this talk we will give a self-contained introduction to performance estimation, describe some of its achievements including a few recent results, and discuss some open questions.

Short bio. François Glineur earned dual engineering degrees from CentraleSupélec and Université de Mons in 1997, and a PhD in Applied Sciences from the latter institution in 2001. After visiting Delft University of Technology and working as a post-doctoral researcher at McMaster University, he joined Université catholique de Louvain, where he is currently a full professor of applied mathematics in the Engineering School. He is also a member of the Center for Operations Research and Econometrics and the Institute of Information and Communication Technologies, Electronics and Applied Mathematics. He served as the Engineering School's vice-dean between 2016 and 2020. His primary areas of interest in research are optimization models and methods, with a particular emphasis on convex optimization and algorithmic efficiency. He received the 2017 *Optimization Letters* best paper award for the worst-case analysis of gradient descent with line search. He is also interested in the use of optimization in engineering, as well as nonnegative matrix factorization and its application to data analysis.

Nicole Megow

University of Bremen

RESOURCE ALLOCATION WITH PREDICTIONS: BALANCING CONSISTENCY AND ROBUSTNESS

Abstract. Expert knowledge, historical data, and machine-learned predictions can guide resource allocation, but relying on them can lead to poor decisions when they are inaccurate or misleading. This talk introduces the framework of algorithms with predictions and explores how to benefit from imperfect information while preserving rigorous performance guarantees. We focus on two settings: multi-stage online bipartite matching, where allocation decisions must be made stage by stage before future demand is revealed, and robust scheduling, where a single schedule must perform well both for a predicted scenario and in the worst case over an uncertainty set. We discuss tradeoffs between consistency with predictions and robustness to uncertainty, highlighting both the possibilities and limitations of combining these guarantees.

Short bio. Nicole Megow is a Professor for Combinatorial Optimization at the University of Bremen. She received her PhD from TU Berlin in 2006 and held research positions at the Max Planck Institute for Informatics, TU Berlin, TU Darmstadt and TU Munich before joining Bremen. Her research focuses on the design and analysis of efficient algorithms for discrete optimization problems, with a particular emphasis on optimization under uncertainty.

Contributed talks

Steven Miltenburg

VU / *On the complexity of capacitated vehicle routing with order restrictions*

The capacitated vehicle routing problem (c -VRP), which is also known as the k -tour cover problem, has been widely studied in many settings. In the Euclidean plane, a PTAS for constant capacity c was given more than 40 years ago, which has been improved and extended ever since. Here, we present arguably the easiest form of a PTAS possible. The algorithm simply chooses between greedy or complete enumeration. A direct corollary of this simplification is that the PTAS extends to basically any kind of order restriction, as long as the objective is to minimize total tour length.

Furthermore, we show that when the metric is a tree, there is a clear separation in complexity between the c -VRP with and without order restrictions. We show that even for a fixed order, the c -VRP is APX-hard for arbitrary (non-constant) capacity c . This in contrast to the polynomial time approximation scheme for the standard c -VRP on trees.

Koen Ligthart

TU/e / *Periodic convexity of integer program value functions and applications to block-structured integer programs*

In this talk, we consider the value functions of integer programs which are of the form $b \mapsto \min\{cx \mid Ax = b, l \leq x \leq u, x \text{ integral}\}$. We argue that such functions exhibit periodic convex behavior that is inherently related to the “integer decomposition property” of polyhedra. This connection can be applied to derive exact algorithms for solving well-structured integer programs. In particular, we settle the parameterized complexity of solving so-called “4-block integer programs”, which are integer programs that decompose into small independent subproblems after removing a small number of variables and constraints. This talk is based on two recent works, joint with Alexandra Lassota and Martin Koutecký.

Yingzhu Liu

RUG / Distributionally Robust Federated Learning with Multi-Source Data

Federated learning trains a shared model from private client data. In practice, data-generating distributions may differ, and the true mixture across clients is often unknown, making the underlying group distribution difficult to specify. Existing approaches address cross-client mixture uncertainty by optimizing against the worst-case mixture, yet assume accurate client-wise distribution estimates. However, these estimates can be unreliable when based on finite samples. To handle both cross-client mixture uncertainty and within-client distributional ambiguity, we construct a global ambiguity set as the union of admissible mixtures of local ambiguity sets. The construction allows client-specific ambiguity radii and admits a client-wise separable reformulation. Leveraging this structure, we establish a high-probability out-of-sample performance guarantee. We further develop a federated algorithm for a penalty-based reformulation and prove its convergence under milder regularity conditions. Simulations validate the algorithm's effectiveness.

Wenhao Chi

UvA / Violation Certification for Linear Predict-Then-Optimize Pipelines

Data-driven decision pipelines combining predictive machine learning models with downstream optimization software are increasingly used to make high-stakes operational decisions. Certifying the safety, fairness, and reliability of these decisions is essential, yet traditional scenario generation methods rely on repeated random testing, which becomes computationally prohibitive when failure events are rare and offers little insight into why failures occur. We present a framework for local violation certification designed specifically for linear decision pipelines under input uncertainty. We mathematically demonstrate that standard sampling methods fail efficiently for rare violations, motivating a direct structural approach. By analyzing the fixed decision boundary of a deployed pipeline, we show that the local risk of failure can be calculated directly in closed form using a single optimization solve. Furthermore, we introduce an exact sampling procedure and closed-form risk statistics that provide feature-level attributions (identifying which input characteristics contribute most to potential non-compliance) without requiring repetitive random trials or complex sampling algorithms. We demonstrate our approach on an economic power dispatch system subject to emissions regulations, delivering precise, auditable risk assessments at a fraction of the traditional computational cost.

Mathijs Barkel

TiU / Integer semidefinite programming for the maximum k -colorable subgraph problem

We study the maximum k -colorable subgraph (MkCS) problem, which consists in finding a largest k -colorable induced subgraph in a given graph. We consider a Semidefinite Programming (SDP) relaxation for the problem and derive two novel families of valid inequalities to strengthen this relaxation. We efficiently solve the strengthened SDP relaxation using a cutting-plane algorithm that is based on the Alternating Direction Method of Multipliers (ADMM). Extensive computational experiments show that the obtained upper bounds outperform the best upper bounds from the literature. To complement our SDP-based upper bounds, we propose an integer ADMM variant that uses an exact Binary Semidefinite Programming (BSDP) formulation of the MkCS problem to produce high-quality feasible solutions. To the best of our knowledge, this is the first application of the ADMM to compute integer solutions to a BSDP problem. We are currently incorporating our methods into a branch-and-bound framework to obtain optimal solutions.

Gijs Bakker

RUG / A dynamic overbooking model for multi-compartment parcel locker systems

We consider the problem of finding an optimal overbooking policy for a single parcel locker with different compartment sizes. Requests for compartments arrive randomly over the course of a period and require an immediate accept or reject decision, without future knowledge on the available capacity or arrivals. Due to the stochastic nature of parcel pickups, optimal control depends not only on revenues but also on pickup rates. In addition, the decision-maker may reallocate uncollected parcels to different compartments during replenishment, referred to as swapping. For a special case, we prove that acceptance decisions are nested in both revenues and pickup rates. For moderate instances, we obtain an optimal policy by dynamic programming and use it to benchmark a proximal policy optimization (PPO) framework, which we find learns a near-optimal policy and scales to larger instances. Lastly, we use a simplified dynamic programming model as a fast heuristic for obtaining a near-optimal locker configuration. Our results imply that acceptance and design decisions should account for pickup times and not only revenues, whereas swapping only marginally improves the objective by at most 3%.

Lucas Ketels

RUG / Convergence Rates for Stochastic Proximal Methods via Regularized Optimality Gaps

We provide a novel analysis of the stochastic proximal point (SPP) algorithm under remarkably weak assumptions. Namely, we provide explicit convergence rates without requiring the objective function to satisfy a regularity condition, such as Lipschitz continuity, which is a very standard assumption in the analysis of this algorithm, or smoothness, less commonly used in the SPP literature. Existing results under our working assumptions provide convergence guarantees, but typically do not provide explicit convergence rates, and the present work fills this gap. Our analysis is based on a non-standard choice of metric for monitoring the progress of the algorithm: rather than the usual optimality gap in the convex setting, we consider a regularized version of the optimality gap. We extend our analysis to a stochastic forward-backward algorithm, where both the proximal backward step and the (sub-)gradient forward step are stochastic. Finally, we emphasize that our analysis does not rely on a strong variance assumption, such as a uniform bound on the gradient variances; instead, we require the variance to be finite only at a solution of the problem.

Daan Noordenbos

TiU / Personalised versus Posted Pricing from Samples

Personalised pricing maximises revenue by tailoring prices to individual customers, but requires detailed information about their willingness to pay. How much of this revenue can be recovered using a posted price based only on samples from their value distribution?

We answer this question by solving a maximin problem: first choosing a pricing rule and then considering its worst-case revenue ratio over the class of λ -regular value distributions.

For the class of monotone hazard rate (MHR) distributions ($\lambda = 0$), posting the sample mean achieves the optimal guarantee $(n/(n+1))^{n+1}$ among all randomised pricing rules. We establish this result by formulating the problem as a doubly infinite linear program.

Beyond the MHR class ($\lambda \in (0, 1]$), optimal sample-based pricing rules become substantially more intricate. Nevertheless, we give a remarkably simple order-statistic-based pricing rule whose performance is within $\Theta(1/n)$ (the optimal rate) of the full-information benchmark $(1 - \lambda)^{1/\lambda}$. Our analysis combines combinatorial identities with techniques from reliability theory.

Wick Wijnholds

UT / Optimal Antenna Design: An Exact Approach

In this talk, we present a new technique for the automated design of the shape of antennas, subject to constraints on their performance. By using mixed-integer programming, we can provide guarantees that none of the state-of-the-art (heuristic) methods can provide. We showcase that even such a non-obvious application of MIPs is feasible once one applies the right transformations to the problem. In particular, we showcase a transformation similar to the Charnes–Cooper transform for linear fractional programming, but applied to linear programs over the complex numbers.

With thanks to our sponsors

NWO and the University of Groningen



**university of
groningen**