Book of abstracts

PGMO DAYS 2026



November 17th and 18th, 2026

EDF Labs Paris-Saclay

Preface

This volume contains the extended abstracts of the talks presented at the conference PGMODAYS 2026 held on November 17th–18th, 2026 at EDF Labs Paris-Saclay.

We especially acknowledge the support of EDF and FMJH. We thank CNRS, Institut polytechnique de Paris, Université Paris-Saclay and Inria. We also thank ROADEF, SMAI-MODE, and SMF. We are grateful to the plenary speakers and to the organizers of the invited sessions.

November 6, 2026
Palaiseau

The organizers of the PGMO Days

Organization Committee

Scientific organizers

Wim van Ackooij	EDF
Claudia d'Ambrosio	CNRS LIX, École polytechnique
Catherine Bonnet	Inria and LSS/Centrale-Supélec
Antonin Chambolle	CEREMADE, CNRS, PSL
Sourour Elloumi	UMA, Ensta Paris and CEDRIC, Cnam (coordinator of PGMO)
Olivier Fercoq	Télécom Paris
Stéphane Gaubert	INRIA and CMAP, École Polytechnique (deputy coordinator of PGMO)
Vianney Perchet	ENSAE Paris (scientific committee chair)

Local organization

Magali Le Chaponnier	Fondation Mathématique Jacques Hadamard
Nadège Menasse	EDF
Mélanie Mihoc	EDF

Scientific Program Committee

Grégoire Allaire	CMAP, Ècole Polytechnique
Christian Artigues	LAAS, CNRS Toulouse
Walid Ben Ameur	SAMOVAR, Télécom SudParis
Claire Boyer	LMO, Université Paris-Saclay
Luce Brotcorne	INOCS, INRIA Lille
Guillaume Carlier	CEREMADE, Université Paris-Dauphine
Antonin Chambolle	CEREMADE, CNRS, PSL
Johanne Cohen	CNRS, LISN, Université Paris-Saclay
Michel De Lara	ENPC, CERMICS
Sophie Demassey	CMA, MINES ParisTech
Frédéric Jean	UMA, ENSTA
Safia Kedad-Sidhoum	CEDRIC, CNAM
Franck Iutzeler	IMT Université de Toulouse
Edouard Pauwels	Toulouse School of Economics
Vianney Perchet	CREST, ENSAE Paris, Chairman
Olivier Spanjaard	LIP6 CNRS, Sorbonne Université

Organizers of Invited Sessions

Andrew McRae:	Nonconvex landscapes and algorithms in data science
Pierre Lavigne:	Recent Advances in Mean Field Games and Mean Field Control
Leo Liberti:	Modelling supply chains in Mathematical Programming: flow and economics
Claudia D'Ambrosio:	Black-box Optimization
Manu Upadhyaya:	Recent Advances in Systematic and Computer-Aided Methods in Optimization
Teemu Pennanen:	Optimal trading in intraday electricity markets
Olivier Bokanowski and Carlos Esteve-Yagüe:	Neural Network approximations for high-dimensional PDEs Problems
Dimitri Papadimitriou:	Quantum Applications in Operations Research
Nicola Paddeu:	Linearly Controlled Systems
Claudia D'Ambrosio:	Mixed Integer Programming and Graphs for Challenging Problems
Luce Brotcorne:	Applications of optimization to power systems and networks
Safia Kedad-Sidhoum:	Discrete optimization for hydro unit commitment problem
Michele Motta:	Geometric Control and Optimization

PGMO 2026 PhD Prizes

The Gaspard Monge Program for Optimization, Operations Research, and their Interactions with Data Sciences awards every year, under the scientific patronage of ROADEF and SMAI-MODE, two PhD prizes (ex aequo). All fields in Optimization and Operations Research, including their interfaces, are eligible. The applicants must have defended their PhD in France during the previous civil year.

The two 2026 PGMO PhD prizes were awarded to

- **Luca Brunod Indrigo**, for his PhD at Sorbonne University and EDF, under the supervision of Pascale Bendotti, Philippe Chrétienne, and Bruno Escoffier, on "Anchor robustness in project scheduling".
- **Romain Cosson**, for his PhD at the Ècole Normale Supérieure – ENS/PSL, under the supervision of Laurent Massoulié, on "Algorithms for Collective Exploration and Decision-Making".

The 2026 PhD prize committee was chaired by Welington De Oliveira (CMA, Mines Paris–PSL). It was composed of the following researchers:

- Emilie Kaufman (CNRS CRISTAL Lille)
- Charles Bertucci (CEREMADE Paris)
- Xavier Venel (LUISS Guido Carli Rome)
- Bernard Penz (GSCOP Grenoble)
- Axel Parmentier (CERMICS, Marne-la-Vallée)
- Sandra U. Ngueveu (LAAS Toulouse)
- Emilie Chouzenoux (INRIA Saclay)
- Adrien Taylor (INRIA Paris)
- Welington De Oliveira (CMA, Mines Paris–PSL)

Plenary Speakers

Christoph Buchheim	Technische Universität Dortmund, Germany — *Combinatorial Optimization in Function Space*
Bissan Ghaddar	IE University and Ivey Business School, Spain — *Sparse Relaxations and Learned Structure for Large-Scale Polynomial Optimization*
Lars Grüne	Universität Bayreuth, Germany — *Performance of Deterministic and Stochastic Model Predictive Control*
Franck Iutzeler	Université de Toulouse, France — *Around distributional robustness and generalization in stochastic optimization*

Table of Contents

Plenary lectures

Combinatorial Optimization in Function Space . . . 1
Christoph Buchheim

Sparse Relaxations and Learned Structure for Large-Scale Polynomial Optimization . . . 2
Bissan Ghaddar

Performance of Deterministic and Stochastic Model Predictive Control . . . 3
Lars Grüne

Around distributional robustness and generalization in stochastic optimization . . . 4
Franck Iutzeler

PGMO PhD Prize lectures

Anchor Robustness and Resource Leveling in Project Scheduling . . . 5
Luca Brunod Indrigo

Collective Algorithms for Exploration and Decision-Making . . . 6
Romain Cosson

Invited and contributed talks

Fixed-Qubit CVaR-VQE and Quantum-Inspired QUBO Sampling for Static and Dynamic Restricted Vertex Multicut in Telecom Resilience . . . 7
Ali Abbassi, Yann Dujardin, Eric Gourdin, Philippe Lacomme, Caroline Prodhon

Probability Maximization via Marginal Assesment of Joint Uncertainties . . . 8
Wim van Ackooij, Carolina Chiu, Welington de Oliveira, Pedro Pérez-Aros

Matrix-based Scenario Aggregation for Two-Stage Stochastic Programs: From Partitions to Subspaces 9
Wim van Ackooij, Natalia Jorquera-Bravo, Pierre-Louis Poirion, Marius Roland

Stochastic Inflows and Seasonal Water Value at Manantali . . . 10
Steeven B. Affognon, Babacar M. Ndiaye, Pierre Mendy, Cheikh M. F. Kebe, Sandrine Charousset, Omar Goudiaby, Thomas Ouillon

Approximate Controllability on the Group of Volume-preserving Diffeomorphisms . . . 11
Andrei Agrachev[1], Bettina Kazandjian[2]

Second-order optimality conditions in geometric optimal control . . . 12
Andrei A. Agrachev, Michele Motta2.

Combining Optimal LEZ Scheduling and EV Purchase Incentives for Transport Policy Backcasting 13
Asmae Alami[1], Vinith Lakshmanan[1], Antonio Sciarretta[1]

Comparing Linear Programming and Mean-Payoff Games via Oriented Matroids . . . 14
Xavier Allamigeon[1], Yiyuan Chen[1], Stéphane Gaubert[1], Xavier Goaoc[2]

Optimization-Oriented Surrogate Frameworks for Wind Farm Layout Optimization . . . 15
Ekaterina ANDREEVA, Claudia D'Ambrosio, Ahlam MAGNOUCHE

Towards Dynamic and Adaptive Tariffs in Electricity Distribution Networks . . . 16
Panagiotis Andrianesis[1], Vasileios Cheimonidis[1], Alexandre Manso de Zuniga[1]

Single-Level Reformulations for Network Pricing under Wardrop Equilibrium with Separable Affine Costs . . . 17
Miguel F. Anjos, José Neto

Online Algorithm Selection: from Foundations to Applications . . . 18
Denis Antipov[1], Carola Doerr[1], Diederick Vermetten[1]

Deterrence of Manipulation through Preference Elicitation in Personalized Incentive Policies . . . 19
Andrea Araldo, Alessa Mayer, Ravi Seshadri, Amir Brudner

Benders decomposition with cut selection for long-term generation and storage capacity expansion problems . . . 20
Ayse N. Arslan, Olivier Baude, Boris Detienne, Aurélien Froger, Rodolphe Griset, Cécile Rottner,

Marius Roland, Frédéric Semet, Mariam Sangaré

Through the G-Polymatroid Glass and What Unit Commitment Finds There 21
Hélène Arvis[1,2,3], Olivier Beaude[1], Nicolas Gast[2], Stéphane Gaubert[3], Bruno Gaujal[2]

Nodal-Aware Flow-Based Capacity Calculation Through Bilevel Inverse Optimization 22
Pierre Arvy, Florin Capitanescu, Amélie Lambert, Daniel Porumbel, Nicolas Omont

Survival of Dominated Strategies under Standard No-Regret Dynamics 23
Galit Ashkenazi-Golan, Domenico Mergoni Cecchelli, Panayotis Mertikopoulos, Edward Plumb, Bassel Tarbush, Yannick Viossat

Benchmarking Deep Learning Architectures for Industrial Non-Intrusive Load Monitoring: Performance, Training Efficiency, and Granularity 24
Laiba Asmatullah, Valentin Colliard, Alizée Mesnard

Using statistical learning in Operations Research: Generalization Bounds and Diffusion Models . . 25
Pierre-Cyril Aubin-Frankowski

Computing Projected solutions of Nash games : from theory to electricity markets 26
Didier Aussel, Jauny, Asrifa Sultana, Shivani Valecha

Stronger Decomposition Cuts for Stochastic Unit Commitment Using Interval Variables 27
Mathis Azema, Vincent Leclère, Wim van Ackooij

Stochastic analysis of an emergency service with two levels sharing the servers 28
Guillaume Ballif[1], Christine Fricker[1]

Mean-Field Generalisation Bounds for Learning Controls in Stochastic Environments 29
Boris Baros[*], Samuel Cohen[*], Christoph Reisinger[*]

Generating electricity with a kite – optimization and bifurcation analysis 30
Antonin Bavoil[1], Romain Veltz[2]

Parameterized Complexity for Segment Routing 31
Cristina Bazgan, Morgan Chopin, André Nichterlein, Camille Richer

Branch-and-Cut en 2 phases pour gérer les symétries polyédrales, application au 1-HUC 32
Pascale Bendotti, Alexandre Heintzmann, Cécile Rottner

A Fresh Look at Lamarckian Evolution and the Baldwin Effect 33
Inès Benito, Johannes F. Lutzeyer, Benjamin Doerr

Lagrangian heuristic for a food lot-sizing problem 34
Lucas Berterottière[1], Adeline Fermanian[2], Safia Kedad-Sidhoum[3]

Pricing for Deliveries Pooling 35
Nicolas Bessin, Guillaume Dalle, Massil Hihat, Frédéric Meunier

Consensus-Based Optimization Beyond Finite-Time Analysis 36
Pascal Bianchi, Radu-Alexandru Dragomir, Victor Priser

Solving Hamilton-Jacobi equations by residual minimization of monotone finite-difference discretizations 37
Olivier Bokanowski, Carlos Esteve-Yagüe, Richard Tsai

Neural Feedback Learning for Stochastic Control with Degenerate Diffusions: Quantitative Error Estimates and Numerical Experiments 38
Olivier Bokanowski, Jean-François Chassagneux, Marco Scaratti, Xavier Warin

The Adjoint State Method for Parametric Definable Optimization without Smoothness or Uniqueness 39
Jérôme Bolte[1], Edouard Pauwels[1], Cheik Traoré[1]

Convergence of Optimizers Implies Eigenvalues Filtering at Equilibrium 40
Jérôme Bolte, Quoc-Tung Le, Edouard Pauwels

An alternative exploration of the solutions space for the Travelling Salesman Problem leveraging a warm-start and a preserving mixer to extend a recent proposal 41
Martin Bombardelli, Philippe Lacomme

Quadratically Regularised Wasserstein Medians 42
Benoît Bonnet-Weill, Matteo Tacchi-Bénard, Nicola Zaupa

Integrated Model for Optimization and Machine Learning: the Case of Predictive Maintenance . . 43
Marie Bouilloud[1], Michel Nakhla[1]

Distributed Optimization for Hydro-thermal Scheduling using ADMM Framework 44
Matias Bresciano, Hamid Ben Ahmed, Roman Le Goff Latimier, Jean-Marc Janin, Juliette Gerbaux

Joint Design of Time-of-Use Schedules and Price Levels: a Bilevel Model with Client Participation and Load Flexibility 45
Abdellah Bulaich Mehamdi, Wim Van Ackooij, Luce Brotcorne, Stephane Gaubert, Quentin Jacquet

A New Bi-Objective View of Bundle Methods 46
Enrico Calandrini[1], Antonio Frangioni[2]

BundleSolver 2.0: A Modular Implementation of Generalized Bundle Methods 47
Enrico Calandrini[1], Donato Meoli[2], Antonio Frangioni[2]

A mathematical programming formulation for generalized supply chains 48
Justin Capgras, Haitam Chakori, Claudia D'Ambrosio, Michele Fabi, Leo Liberti

A mathematical programming formulation for generalized supply chains - Industrial and Economic mechanisms 49
Justin Capgras, Haitam Chakori, Claudia D'Ambrosio, Michele Fabi, Leo Liberti

Trading with market resistance and concave price impact 50
Nathan De Carvalho[1], Youssef Ouazzani Chahdi[2], Grégoire Szymanski[3]

Optimal Power flow in Transmission Networks: a Spectral Graph Theory approach 51
G. Cavalcante, L. Melo, A. Knippel, J.-G. Caputo, R. Andrade

Black-box adversarial attacks on SPD Networks 52
Bastien Cavarretta, Thibault de Surrel, Clément W. Royer, Florian Yger

Learning controls in mean games and mean field control problems 53
Meynard Charles

Numerical bounds for European option prices in uncertainty volatility model. 54
Jean-François CHASSAGNEUX[1], Simon SANANES[2]

Geometric Routing in Geometric Inhomogeneous Random Graphs 55
Yu-Cheng Chiu, Marc Kaufmann, Kostas Lakis, Ulysse Schaller

ProbLiM: A Julia Toolbox for Linear Problems with Chance Constraints 56
Yi Chu, Welington de Oliveira, Wim van Ackooij

On convergence rates of subgradient descent on semialgebraic functions 57
Evgenii Chzhen, Sholom Schechtman

Deep Hilbert–Galerkin Methods for Infinite-Dimensional PDEs and Optimal Control 58
Samuel N. Cohen[1], Filippo de Feo[2], Jackson Hebner[1], Justin Sirignano[1]

On Stochastic Nonconvex Bilevel Optimization 59
Daniel Cortild[1], Mathias Staudigl[2], Juan Peypouquet[3], Coralia Cartis[1]

Decision-Focused Learning for Dynamic Inventory Replenishment under Production Constraints . . 60
Nicolas Corvol[1], Louis Bouvier[1], Axel Parmentier[1]

Optimizing Muli-Year Investment Planning for Electricity Distribution Grid Resilience 61
Axelle Cousinou[1], Sebastien Delbos[1], Mathieu Francois[1]

Scheduling of Nuclear Power Plant Outages under Duration Uncertainty 62
Ezra Daniels, Miguel F. Anjos, Luce Brotcorne, Bárbara Rodrigues, Marius Roland

Market Equilibrium based Co-Opetitive Infrastructure Sharing 63
Mandar Datar, Andrea Araldo, Tijani Chahed

Decision-Focused Policies for Microgrid Energy Management 64
Alexis David[1], Tristan Rigaut[1], Vincent Leclère[2], Adrien Le Franc[2]

From Random Games to Deterministic Limits: A Differential-Game Approach to Homogenization 65
Andrea Davini, Raimundo Saona, Bruno Ziliotto

A Fluid Optimization Approach for Decomposable MDPs with Combinatorial Actions: Application to Offshore Substation Maintenance . . . 66
Solène Delannoy-Pavy, Axel Parmentier, Vincent Leclère, Manuel Ruiz, Cyrille Vessaire

Regularization by Common Noise for Mean Field Games and Mean Field Control . . . 67
François Delarue, Lukas Wessels

Surrogate-Based Optimization for Adversarial Sonar Placement . . . 68
Lorenzo Demaria, Axel Kärrolm, Jan Kronqvist

Scheduling for "Ramasse": How to Collect Goods on the Rungis International Market . . . 69
Natello Descormier, Lucas Berterottière, Pierre Fouilhoux, Massil Hihat, Safia Kedad-Sidhoum

On the notion of Wasserstein BSDE and its applications . . . 70
Mao Fabrice Djete[1], Mattia Martini[2]

Understanding the Co-evolutionary Solution Exchange of the MOEA/D: Minimal Communication Gives Maximal Speed-Ups . . . 71
Benjamin Doerr, Moaad Elmoutassim, Martin S. Krejca, Guy Mushiata Kibulu

Full-Dimensional Many-Objective Benchmarks for Discrete Multi-Objective Optimization . . . 72
Benjamin Doerr, Martin S. Krejca, Dieks Scholten

Graph powers for robust optimization . . . 73
Alexandre Dupont-Bouillard, Leo Liberti, Antonio Mucherino

Pricing for Real-Time Ride-Sharing with Random Demand and Supply . . . 74
Rose Sossou Edou, Virginie Boutueil, Frédéric Meunier

On the Explainability of Max-Plus Neural Networks . . . 75
Ikhlas Enaieh, Olivier Fercoq, Angel Garcia Pedrero

Benign Landscape Of Nonconvex Low-dimensional Relaxation Of Orthogonal Synchronization With Group-symmetric Dual Certificates . . . 76
Faniriana Rakoto Endor[1], Irène Waldspurger[2]

Optimal control under uncertainty for relative guidance during orbital rendezvous . . . 77
Laura Bernardo Estévez, Prince Edorh, Bruno Hérissé, Frédéric Jean

An AMM for Energy Sharing: From Theory to Verifiable Settlement . . . 78
Michele Fabi[1], Baptiste Leloup[2]

A mixer for minimization of function with constraint . . . 79
Gerard Fleury, Philippe Lacomme

Specification of the GDEMO-II Variant for Box Archiving . . . 80
Jules Fraizier, Benjamin Doerr, Martin S. Krejca

Some curiosities about the convergence of the Moment-SOS hierarchy . . . 81
Adrien Le Franc

Exact Reformulation for Nonconvex Wasserstein DRO: Application to Day-Ahead Renewable Nomination . . . 82
Charles Garrisi, Welington de Oliveira, François Pacaud

Adjoint-based Optimality Conditions for a Probability Maximization Problem . . . 83
Caroline Geiersbach, Wim van Ackooij, Hasnaa Zidani

Combinatorial Optimization for Energy-Efficient In-Network Learning . . . 84
Garance Gérard, Mustapha Bouhtou, Safia Kedad-Sidhoum, Amélie Lambert, Nancy Perrot

Graph Partitioning and ADMM for Optimal Pump Scheduling in Drinking Water Distribution Networks . . . 85
Carlo A. M. Gervasi, Stefan Ropke, Claudia D'Ambrosio, Bissan Ghaddar

Online Combinatorial Learning with Causal Rescheduling for EV Charging . . . 86
Eymeric Giabicani[1], Raphaël Féraud[4], Patrick Maillé[2], Guy Camilleri[3], Anne Blavette[1]

Limited-Memory First-Order Methods and H -Duality 87
Pierre Vernimmen[1], François Glineur[1]

Generalized infinite dimensional Alpha-Procrustes based geometries 88
Salvish Goomanee[*†], Andi Han[‡], Pratik Jawanpuria[§], Bamdev Mishra[¶]

Topology optimization of PEMFC channels using a deshomogenization method 89
Luca Gorini[1], Sébastien Kawka[1]

A Branch&Check approach for Selecting Multi-Site Balancing Entities 90
Emilien Gouffault, Alexandre Marié, Emily Little

Trust-Region Methods For Burer-Monteiro Factorizations Of Phase Retrieval Problems 91
Florentin Goyens, Clément W. Royer, Irène Waldspurger

Aperiodic Tiling Rhythmic Canons 92
Gennaro Auricchio[1], Luca Ferrarini[2], Camille Grange[2]

EvoHGS: Co-evolving Language Models with Hybrid Genetic Search for Vehicle Routing Problems . 93
Zihao Guo, Yiding Shi, Xia Jiang, Yaoxin Wu, Jie Gao, Andrea Araldo, Mounîm El-Yacoubi

Parametric Physics Informed Neural Network for the Finite Horizon classical Lady-in-the-Lake Pursuit-Evasion Differential Game 94
Dhruv Gupta

Superposition Principles for Perturbed Sweeping Processes and Differential Inclusions 95
Aldo Gutiérrez-Concha, Francisco Silva, Emilio Vilches

Drawing and Scheduling Incomplete Tournaments with Seeding Constraints: the Case of the UEFA Champions League 96
Julien Guyon, Adle Ben Salem, Thomas Buchholtzer, Mathieu Tanré

Coupled Investments in SMRs and Industrial Sites: a Real Options Analysis 97
Cyrian Hallermeyer[1,2,3]

Chance-constrained Linear Complementarity Problems 98
René Henrion, Martin Schmidt

Piecewise linear approximation for the hydro unit commitment problem 99
Kea Horvath, Christian Artigues, Pascale Bendotti, Sandra Ulrich Ngueveu, Cécile Rottner

On the Complexity of some Unicity Problems 100
Olivier Hudry

Landscape Analysis for Shallow Neural Networks: Complete Classification of Critical Points for Cubic Activation and Affine Target Functions 101
Shokhrukh Ibragimov[1], Ilkhom Mukhammadiev[2], Diyora Salimova[3]

Quantum Genetic Algorithm for Combinatorial Optimisation 102
Adil Kassaoui, Eric Bourreau

Convergence Rates for Stochastic Proximal Methods via Regularized Optimality Gaps 103
Lucas Ketels, Daniel Cortild, Juan Peypouquet, Guillaume Garrigos

Variational Quantum Algorithms for the Uncapacitated Single-Allocation p-Hub Median Problem . 104
Rashed Khanjani-Shiraz[1]

Choice-Based Dynamic Pricing with Deferred Departures in Bike-Sharing Networks 105
Mesut Can Köseoğlu, Shadi Sharif Azadeh, Luce Brotcorne

Feasibility, Cost, and Environmental Assessment of Heavy Duty Truck Fleet Electrification via Mixed Integer Optimization 106
Malak Al Kostit[1,2], Sophie Demassey[1], Valentina Sessa[1], Franck Tirard[2], Pierre Chaufour[2]

Beyond Softmax and Entropy: Convergence Rates of Policy Gradients with f-SoftArgmax Parameterization & Coupled Regularization 107
Safwan Labbi[1], Daniil Tiapkin[1,2], Paul Mangold[1], Eric Moulines[3,4]

Target-Volume Formulations in the Hydro Unit Commitment Problem 108
Marguerite Lafeuillade, Sourour Elloumi, Safia Kedad-Sidhoum, Cécile Rottner

Optimization of the energy consumption of buildings . . . 109
Michael Laporte[1], Claudia D'Ambrosio[2], Frederic Meunier[1]

What Makes Information More Valuable? An Answer With Convex Analysis . . . 110
Michel De Lara[1]

A conditional gradient approach to regularized mean field optimization problems . . . 111
Pierre Lavigne[*]

The Moment-SOS Hierarchy: Applications to Reaction-Diffusion PDEs and Navier–Stokes Equations 112
Charlie Lebarbé, Didier Henrion, Milan Korda, Émilien Flayac, Michel Fournié

Path-conditioned training: a principled way to rescale ReLU neural networks . . . 113
Arthur Lebeurrier[1], Titouan Vayer[2], Remi Gribonval[3]

A Clustered Team Orienteering Problem with Synchronization . . . 114
Aymeric Legros[1,2], Alain Hertz[2]

Linear Analytical Model for Static Traffic Assignment: Explicit Formulation and Optimization-Oriented Applications . . . 115
Baptiste Leroux[1], Antonio Sciarretta[1], Milos Balac[2]

Convergence of empirical subgradients for optimal transport-based objectives . . . 116
Tam Le[1]

Statistics and optimization in intraday electricity markets . . . 117
Tanguy Lindegaard

A Complex-Valued Continuous-Variable Quantum Approximation Optimization Algorithm (CCV-QAOA) . . . 118
Raneem Madani, Abdel Lisser, Zeno Toffano

Robust stochastic optimization with risk-averse constraints via variance penalization and operator splitting. . . . 119
Paul Malisani, Adrien Spagnol, Fabrice Guillemin

Decomposition of Convex Large Scale Stochastic Optimization Problems . . . 120
Paul Malisani, Laurent Pfeiffer, François Caffier

Data-Driven Security Assessment of Distribution Grids with Weibull-Based Synthetic Data Generation 121
Alban Marie[1], Juan J. Cuenca[2], Emanuel Aldea[3], Guy Camilleri[4], Anne Blavette[1]

On ℓ_0-Regularized Sparse Multiple Kernel Learning for Feature Selection . . . 122
Rui Ma[1], Rodolphe Le Riche[1]

Attention-based PCA . . . 123
Rodrigo Maulen-Soto[♠], Claire Boyer[♢♣]

opt-syn: Analysis and Synthesis of First-Order Algorithms . . . 124
Jared Miller[1], Manuel Zobel[1], Fabian Jakob[2], Andrea Iannelli[2], Carsten Scherer[1]

Optimizing Hospital Procurement Decisions: The Case of a University Hospital . . . 125
Cecília Mota, Lara Pontes, Luciano Costa, Anand Subramanian

Multimodal Optimization of K-means Clustering . . . 126
Ravil Mussabayev, Rustam Mussabayev

Discovering Interpretable Multi-Parameter Control Policies for Evolutionary Algorithms Using Deep Reinforcement Learning . . . 127
Tai Nguyen, Phong Le, Carola Doerr, Nguyen Dang

Reinforcement Learning for Solving Multimodal Functions with CMA-ES . . . 128
Marc Oedingen, Oskar Girardin, Nikolaus Hansen, Alexandre Chotard

Rank-Adaptive Tensorized Neural Operators for PDE-Constrained Optimization . . . 129
Younes Brahim Oumedjber[1], Akram Boukabache[2], Nadjette Debbouche[2], Adel Ouannas[3]

Interior-point methods for large-scale equilibrium problems . . . 130
François Pacaud, Benoît Legat, Quentin Lété, Quentin Jacquet

Quantum Computing for Integer Quadratic Programming . . . 131
Dimitri Papadimitriou

Convex optimization in intraday electricity markets . . . 132
Teemu Pennanen

Differentiating Minimal-Norm Solutions to Parametric Optimization Problems . . . 133
Baptiste Plaquevent-Jourdain[1], Jalal Fadili[2], Antonio Silveti-Falls[1]

Probing-Enhanced Stochastic Programming: Structural Dominance, Stability, and Approximate Branch-and-Bound . . . 134
Franco Quezada, Céline Gicquel, Safia Kedad-Sidhoum, Bernardo Pagnoncelli

A New Solution Concept for Games Allowing for Infeasibility . . . 135
Anne Reulke[1,2], Mikaël Touati[1], Rachid El-Azouzi[2]

Robotaxi Dispatching with Endogenous En-route Times . . . 136
David Rey[1]

Global Convergence of Second-Order Methods for Physics-Informed Neural Networks . . . 137
Konstantin Riedl[1], Justin Sirignano[2], Konstantinos Spiliopoulos[3]

Large-Scale Analysis of Multi-Scale Queuing Networks for Car-Sharing Systems . . . 138
Alessia Rigonat, Christine Fricker

Competitive Menu Pricing for Electric Vehicle Charging Stations . . . 139
Rita Safi, Margarida Carvalho, Yezekael Hayel, Raphaël Payen, Julien Pennec, Tania Jiménez

Wasserstein-Fisher-Rao differential calculus and optimal control of weighted particle systems . . . 140
Enrico Sartor[1]

Spectral-Gauge Cuts for Semidefinite Optimization . . . 141
Antonio Sasaki[1], Sophie Demassey[1], Valentina Sessa[1]

Solving Minimum-Cost Flows on GPUs . . . 142
Simeon Schaub, Guillaume Dalle

SS-TEA: Search-Supervised Transformer Edge Addition for Black-Box Graph Intervention . . . 143
Xu Shuangshuang (Mme)[1], Andrea Araldo[2], Duo WANG[3]

Exact Extremal Optimization for Wrong-Sign Entries in Equal-Diagonal Symmetric Matrices . . . 144
Harman Singh

Monotone Inclusions for Decision-Dependent Stochastic Problems . . . 145
Wutao Si, Hamza Ennaji, Jalal Fadili

Blackbox Optimization for Active Distribution Networks Reconfiguration under Uncertainty . . . 146
Christina G. Soldati, Hussein Suprême, Miguel F. Anjos, Sébastien Le Digabel, Antoine Lesage-Landry

Combinatorial Optimization Augmented Machine Learning for Dynamic Electric Autonomous Dial-a-Ride Problem . . . 147
Yue Su[1], Léa Ricard[2]

Convergence rates for the moment-SoS hierarchy . . . 148
Corbinian Schlosser[1], Matteo Tacchi-Bénard[2], Alexey Lazarev[3]

Branching of Regular Abnormal Geodesics in Sub-Riemannian Manifolds of Rank Two . . . 149
Luca Rizzi[1], Mario Sigalotti[2], Nicolò Tedesco[1,2]

Metric limits of sub-Finsler Lie groups . . . 150
Enrico Le Donne[1], Sebastiano Nicolussi Golo[2], Andrea Tettamanti[1]

Computer-Aided Proofs in First-Order Optimization, with Applications to Error Feedback . . . 151
D. B. Thomsen, A. Taylor, A. Dieuleveut

Simple Proofs, and How to Find Them . . . 152
Daniel Berg Thomsen, Manu Upadhyaya, Baptiste Goujaud, Aymeric Dieuleveut, Adrien Taylor

Sensitivity of GPU Scheduling Rules to Runtime Estimation Error . . . 153
Yosef Gavriel Tirat-Gefen, Filipe Vinci

When does residual minimization solve a differential equation? . . . 154
Richard Tsai, Carlos Esteve-Yagüe, Stanley Osher

Pontryagin Maximum Principle for Non-Smooth State-Constrained Control Problems over Wasserstein Spaces . . . 155
Fernanda Urrea[1], Ernesto Treumún[2]

Sensitivity Analysis for free terminal time horizon problems in Wasserstein spaces . . . 156
Fernanda Urrea[1]

Implicit Bias of SGD for Overparametrized Models . . . 157
Govind Menon[1], Austin J. Stromme[2], Adrien Vacher[2]

Controlled Screening for Smooth Objectives . . . 158
Eyal Vayness[1], Olivier Fercoq[2], Maxime Sangnier[1]

Optimizing Factorable Losses via a Change of Variable: Properties and Computational Opportunities 159
Manon Verbockhaven, Arthur Lebeurrier, Rémi Gribonval

What is the Objective? Investigating Design Decisions in Continuous Multi-Objective Optimization 160
Diederick Vermetten

Can Acceleration in Gradient-Norm Minimization Be Anytime? Sharp Last-Iterate Limits in Smooth Convex Optimization . . . 161
Pierre Vernimmen[1], François Glineur[1]

Primal-Dual Coordinate Descent for Nonconvex-Nonconcave Saddle Point Problems Under the Weak MVI Assumption . . . 162
Iyad Walwil[1], Olivier Fercoq[2]

Online Design of Dynamic Networks . . . 163
Duo Wang, Andrea Araldo, Mounîm El-Yacoubi

Learning PDEs for Portfolio Optimization with Quantum Physics-Informed Neural Networks . . . 164
Letao Wang, Abdel Lisser, Sreejith Sreekumar, Zeno Toffano

Long-time behavior of optimal particle systems in mean-field control . . . 165
Songbo Wang[1], Wenbin Yan[2]

Optimized Projection-Free Algorithms for Online Learning: Construction and Worst-Case Analysis 166
Julien Weibel, Pierre Gaillard, Wouter Koolen, Adrien Taylor

Quantum Approaches to Multi-Objective Optimization . . . 167
Laura Wendel, Alexey Bochkarev, Zakaria Abdelmoiz Dahi, Pascal Halffmann

Opportunities in Algorithmic Trading for Battery Energy Storage Systems in Short-Term Electricity Markets . . . 168
David Wozabal

Profit-Optimized Maintenance Task Clustering for Gas-Insulated Substations . . . 169
Wenbo Wu[1], Rim Kaddah[1], Miguel F. Anjos[2], Sihem Tebbani[3]

Eventual convexity for separable chance constraints with skewed generalized hyperbolic random variables . . . 170
Heng Zhang, Abdel Lisser

Second-Order Reliability Method based on Gram-Charlier A series with analytical Gradient and Hessian of failure probability . . . 171
Qian Zhang, Mårten Olsson, Rami Mansour

A Crowding Distance That Provably Solves the Difficulties of the NSGA-II in Many-Objective Optimization . . . 172
Weijie Zheng, Yan Gao, Benjamin Doerr

Session 14: Plenary 4 : Christoph Buchheim *Room: Marie Curie (Auditorium)*

Combinatorial Optimization in Function Space

Christoph Buchheim

TU Dortmund, Germany, christoph.buchheim@math.tu-dortmund.de

Optimal control problems with discrete-valued controls are typically tackled via a first-discretize-then-optimize strategy, in which the original infinite-dimensional problem is approximated by a large-scale finite-dimensional mixed-integer nonlinear program (MINLP) and handed to a general-purpose solver. A fundamental drawback of this approach is that conventional discretization schemes shrink the feasible set, so that dual bounds on the discretized problem are not valid in general for the original infinite-dimensional formulation. Moreover, such bounds do not carry over when the mesh is refined. Although one can in principle correct for this by estimating the discretization error, the resulting procedure is both technically involved and tends to produce only weak bounds.
To overcome these limitations and to support a branch-and-bound-and-refine scheme for global optimization in function space, we introduce a first-dualize-then-discretize paradigm. In this framework, a convex relaxation of the original problem is dualized directly in function space; the resulting dual problem can then be discretized with any conforming scheme, and the dual objective value of the discretized dual problem constitutes a valid lower bound for the original infinite-dimensional problem. We demonstrate the approach on two classes of problems featuring total variation or minimum dwell-time constraints, for which we also provide extended formulations in function space.

Session 6: Plenary 2 by Bissan Ghaddar *Room: Marie Curie (Auditorium)*

Sparse Relaxations and Learned Structure for Large-Scale Polynomial Optimization

Bissan Ghaddar

IE University and Ivey Business School, Spain, bghaddar@faculty.ie.edu

Polynomial optimization problems (POPs) arise in numerous domains such as energy systems and water networks, yet their nonconvexity poses major computational challenges. This talk focuses on exploiting sparsity structures to enable scalable global optimization of large-scale POPs. We develop formulations that leverage correlative sparsity to construct structured semidefinite and second-order cone relaxations of significantly reduced dimensionality. We further explore machine learning models that learn variable dependencies in order to guide sparsity-pattern selection. By exploiting chordal extensions of the sparsity graph and decomposing the semidefinite matrices, the resulting sparse relaxations achieve tight bounds at tractable computational cost. Computational results on optimal power flow instances demonstrate the effectiveness of the approach.

Session 2: Plenary 1 : Lars Grüne *Room: Marie Curie (Auditorium)*

Performance of Deterministic and Stochastic Model Predictive Control

Lars Grüne

Universität Bayreuth, Germany, lars.gruene@uni-bayreuth.de

Model Predictive Control (MPC) is one of the most successful modern control techniques. It generates a feedback control by successively solving finite-horizon optimal control problems. A key question is whether the resulting feedback is approximately optimal in some sense. In this talk we give a positive answer to this question using dissipativity and turnpike properties. We first describe results for deterministic MPC, then present an industrial use case, and finally give a survey on recent results for stochastic MPC schemes.

Session 11: Plenary 3 : Franck Iutzeler *Room: Marie Curie (Auditorium)*

Around distributional robustness and generalization in stochastic optimization

Franck Iutzeler

Université de Toulouse, France, franck.iutzeler@math.univ-toulouse.fr

Transferring predictive power from past data to similar future situations is a central objective in statistical learning and stochastic optimization. In the past years, distributionally robust optimization has become popular by providing a mathematically elegant proxy for this problem by relying on optimal transport between distributions. We will cover recent advances on this approach but also its limitation in deep learning and related perspectives.

Session 4: PhD Prize Lecture 1 : Luca Brunod Indrigo *Room: Marie Curie (Auditorium)*

Anchor Robustness and Resource Leveling in Project Scheduling

Luca Brunod Indrigo

Sorbonne Université, France, luca.brunod.indrigo@gmail.com

This talk presents Luca Brunod Indrigo's thesis, supervised by Pascale Bendotti, Philippe Chrétienne and Bruno Escoffier, which was defended in December 2025. Scheduling is the field of optimization that aims at deciding when to process the jobs of a project so that certain constraints are respected and a performance criterion is maximized. In a real-world context, a scheduling problem typically needs to be solved at various stages of a project, between which uncertainty may disrupt project data. At each stage, new scheduling decisions are expected to match the decisions that were taken at earlier stages as best as possible. Indeed, deviating too much from previous decisions can entail additional costs, e.g., penalties from subcontractors, and should therefore be avoided. The thesis is organized in two parts. The first part deals with a special case of robustness called anchor robustness that aims at guaranteeing the starting times of some jobs. Several problems are studied that consider both the proactive aspect, i.e., building schedules with some guaranteed starting times, and the reactive aspect, i.e., rescheduling once uncertainty is realized. The second part deals with resource leveling, which in the present case is to respect a given resource level as best as possible. Polynomial solution methods, NP-hardness proofs as well as approximation results are provided for such problems. This thesis is an industrial thesis in cooperation between EDF, the main French power provider, and the LIP6 laboratory at Sorbonne-Université. The research carried out is inspired by problems found in the industrial context of EDF. More precisely, the main motivation for this work is the scheduling of large-scale maintenance tasks performed during outages of power plants. Such projects contain many jobs, involve a variety of subcontractors and must be planned sufficiently in advance, which makes them prone to uncertainty. The heavy equipments and the specialized workforce that take part in those projects are rare resources that need to be properly managed to avoid delays and additional costs. For all these reasons, anchor robustness and resource leveling are helpful as dedicated optimization tools to improve performance and reliability at each stage of planning. This aspect of practical applicability has been a central concern all along and some algorithms have indeed been implemented for life-size instances.

Session 7: PhD Prize Lecture 2 : Romain Cosson *Room: Marie Curie (Auditorium)*

Collective Algorithms for Exploration and Decision-Making

Romain Cosson

New York University, United States of America, rc6142@nyu.edu

This talk presents Romain Cosson's thesis, which was defended in team Argo of Inria Paris in July 2025, under the supervision of Laurent Massoulié, and which was selected for the PGMO PhD award 2026. The thesis focuses on the design and analysis of online algorithms, which make decisions based on information revealed progressively. It is centered on problems in which several decision-makers must coordinate to perform tasks with limited information, such as exploring an unknown structure. A central problem studied in the thesis is collective tree exploration, which models the exploration of an unknown environment by multiple mobile agents (see an illustration here). The thesis also investigates other problems, such as ultrametric traversal (illustration here), metrical task systems, and the k-server problem. The thesis proposes new algorithms, improves performance guarantees, and establishes new lower bounds. It also investigates various aspects of online-algorithm complexity, such as average-case analysis, communication complexity, computational complexity, and randomness.

Session information not available in the current online program.

Fixed-Qubit CVaR-VQE and Quantum-Inspired QUBO Sampling for Static and Dynamic Restricted Vertex Multicut in Telecom Resilience

Ali Abbassi, Yann Dujardin, Eric Gourdin, Philippe Lacomme, Caroline Prodhon

[1]Orange Research, Châtillon, France; [2]LIST3N, Université de technologie de Troyes, France; [3]LIMOS, UMR CNRS 6158, ISIMA, Aubière, France

Keywords: restricted vertex multicut, row generation, QUBO, CVaR-VQE, finite-horizon optimization, telecom resilience

Context and positioning. Isolating network assets so that protected source–sink pairs can no longer communicate leads to the restricted vertex minimum multicut problem. Its compact decision space hides exponentially many path inequalities. We let a classical row-generation layer handle separation, feasibility verification and constraint generation, while quantum or quantum-inspired samplers solve the current restricted problem under an exact certifying reference.

Static and dynamic problems. Let $G = (V, E)$, let H be the set of terminal pairs, let V_H be protected terminals, and let $R = V \setminus V_H$ be removable vertices with costs $c_v \geq 0$. With $x_v = 1$ when $v \in R$ is removed and $\mathcal{P}_{s,t}(G)$ the simple s–t paths, the static problem is

$$\min \sum_{v \in R} c_v x_v \quad \text{s.t.} \quad \sum_{v \in \pi \cap R} x_v \geq 1 \quad \forall (s,t) \in H,\ \forall \pi \in \mathcal{P}_{s,t}(G), \qquad x \in \{0,1\}^R.$$

The finite-horizon extension uses stage-dependent graphs, pairs and costs and adds the reconfiguration penalty

$$\min \sum_{t=1}^{T} \sum_{v \in R_t} c_{v,t} x_{v,t} + \sum_{t=2}^{T} \sum_{v} \rho_v |x_{v,t} - x_{v,t-1}|,$$

subject to stage-wise path inequalities.

Hybrid row generation. A separation round tests reachability in the residual graph and adds surviving paths. Exact restricted solves certify global optimality at termination; heuristic solves certify graph feasibility of the returned cutset only. With $L_\pi = |\pi \cap R|$ and $b_\pi = \lceil \log_2 L_\pi \rceil$ slack bits, the exact-slack QUBO is

$$Q(x, y) = \sum_{v \in R} c_v x_v + M \sum_{\pi \in \mathcal{P}_{\text{act}}} \left(\sum_{v \in \pi \cap R} x_v - 1 - \sum_{q=0}^{b_\pi - 1} 2^q y_{\pi q} \right)^2 .$$

If $M > \max_{v \in R} c_v$, every global minimizer satisfies all active path inequalities and minimizes deletion cost among active-row-feasible assignments. The fixed-register CVaR-VQE alternative keeps one qubit per removable vertex and scores bitstrings by deletion cost plus a penalty for unhit active paths. Generated paths modify the loss but do not enlarge the register. Dynamic switching adds couplings between consecutive stages but no new QUBO variables beyond the stage decisions.

Numerical experiments. The static benchmark contains 425 instances from 22 SNDlib topologies; the dynamic benchmark contains 386 five-stage scenarios. Exact mixed-integer row generation provides the certified optimum. In the static experiments all methods return graph-feasible cutsets. Exact row generation is fastest and certifying; slack-QUBO annealing matches the optimum on 281/425 instances, while the fixed-register CVaR-VQE pipeline matches it on all 425. Sampler-controlled experiments show no replicated gain attributable to the quantum circuit itself. In the dynamic experiment, temporal coupling is the bottleneck: slack QUBO stays feasible on most scenarios but switches often, while CVaR-VQE has lower coverage because of the horizon-wide qubit register and iteration limits.

Discussion. Only the slack encoding grows with the active constraint set, whereas the fixed register grows with the horizon. Rolling windows can reduce the latter at the cost of full-horizon optimality. The experiments establish no quantum computational advantage; rather, they show that any quantum component must be assessed within the complete workflow, including graph validation and temporal coupling.

References

[1] A. Abbassi et al., CoDIT 2025. [2] A. Abbassi et al., CCIS 2744, 2026. [3] P. K. Barkoutsos et al., Quantum 4:256, 2020. [4] M. Cerezo et al., Nature Reviews Physics 3:625–644, 2021. [5] F. Glover, G. Kochenberger, Y. Du, 4OR 17:335–371, 2019. [6] A. Lucas, Frontiers in Physics 2:5, 2014. [7] Y. Magnouche et al., Discrete Applied Mathematics, 2021. [8] D. Marx, I. Razgon, SIAM J. Comput., 2014. [9] S. Orlowski et al., Networks 55(3):276–286, 2010. [10] A. Peruzzo et al., Nature Communications 5:4213, 2014. [11] G. Vidal, Phys. Rev. Lett. 91:147902, 2003.

Session 8E: Optimisation with chance constraints *Room: Antoine Lavoisier (A1.122)*

Probability Maximization via Marginal Assesment of Joint Uncertainties

Wim van Ackooij, Carolina Chiu, Welington de Oliveira, Pedro Pérez-Aros

1OSIRIS, EDF Lab Paris-Saclay, France, wim.van-ackooij@edf.fr 3MINES Paris-PSL, CMA, France, welington.oliveira@minesparis.psl.eu 2Universidad de Chile, Chile, cchiu@dim.uchile.cl, pperez@dim.uchile.cl

Keywords: Probability maximization, stochastic optimization, first-order methods, L-smooth regularity

We consider probability maximization problems of the form

$$\max_{x\in X} \ \phi(x) := \mathbb{P}(\xi \leq h(x)),$$

where $X \subset \mathbb{R}^n$ is a closed feasible set, $h : \mathbb{R}^n \to \mathbb{R}^m$ is a $C^{1,1}$ mapping, and ξ is a random vector with continuously differentiable density, and marginal distributions F_j, $j = 1, \dots, m$. After establishing conditions under which ϕ is a $C^{1,1}$ function, we propose MAJU (Marginal Assessment of Joint Uncertainty), an algorithm that exploits the marginal distributions F_j to construct a local approximation model at the current stability center $\hat{x}^k$:

$$M_k(x) = \phi(\hat{x}^k) + \sum_{j=1}^{m} w_j^k (F_j(h_j(x)) - F_j(h_j(\hat{x}^k))).$$

Here, $w^k \in [0,1]^m$ is a vector of weights selected by the algorithm, which defines the next iterate as a stationary point of the proximal subproblem

$$\max_{x\in X} \ M_k(x) - \frac{L_k}{2}\|x - \hat{x}^k\|^2.$$

A serious-step/no-step acceptance rule, together with the update of $L_k > 0$, promotes stability of the method and ensures monotone ascent of the objective function. We illustrate the performance of MAJU on three benchmark probability maximization problems, where it exhibits competitive performance compared with other known NLP methods applicable to $C^{1,1}$ objectives.

References

[1] W. van Ackooij, C. Chiu, W. de Oliveira, and P. Pérez-Aros, Lipschitz gradient guarantees for probability functions and a new algorithm for probability maximization, preprint, available at Optimization Online, 2026.

Session 5E: Handling Large scale Stochastic Programs *Room: Antoine Lavoisier (A1.122)*

Matrix-based Scenario Aggregation for Two-Stage Stochastic Programs: From Partitions to Subspaces

Wim van Ackooij, Natalia Jorquera-Bravo, Pierre-Louis Poirion, Marius Roland

1EDF Lab Paris-Saclay, France, wim.van-ackooij@edf.fr 2 Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL, France, natalia-francisca.jorquera-bravo@inria.fr, mmmroland@gmail.com 3Center for Advanced Intelligence Project, RIKEN, Japan, pierre-louis.poirion@riken.jp

Keywords: Two-Stage Stochastic Programming, Scenario Aggregation, Matrix Decomposition

In this work, we study two-stage stochastic programs in which uncertainty is represented by a finite set of scenarios. Accurately representing the uncertainty requires a large number of scenarios. This makes the extensive formulation intractable, and has motivated a wide range of scenario reduction, clustering and decomposition techniques [2, 1]. We first propose a matrix-based view of the problem that brings them together. The idea is to aggregate the constraints and the variables of the second stage through two matrices. Each row of a matrix defines an aggregated constraint or an aggregated variable, and its entries give the contribution of each scenario to it. We show that every scenario clustering scheme is a particular choice of these two matrices. Adaptive partition-based methods, which refine a partition of the scenario set guided by dual information, correspond to matrices whose rows have disjoint supports and binary entries. We derive a bound on the approximation error at any first-stage decision. This bound measures the distance between the optimal dual solution and the solutions that the aggregation matrix can reach. For a partition-based method, it is governed by the heterogeneity of dual solutions inside each subset in the partition. In that case we also show that the two matrices can be used together while keeping a relaxation of the original problem, so the partition reduces the number of variables at no cost on the quality of the bound. This property is specific to the partition structure and fixed recourse assumptions, and it gives insights on the advantages of adaptive partition-based methods. This framework leads us to a new algorithm. We keep only the aggregation of the constraints, and we let the rows of the matrix have arbitrary entries and overlapping supports. A constraint can then contribute to several of the new constraints at the same time. The approximation matrix is now built by factorizing the matrix of the dual solutions collected along the iterations. Unlike adaptive partition-based methods, our algorithm does not require the fixed recourse assumption. Finally, we test our algorithms on instances with up to ten thousand scenarios to validate their performance and practical applicability.

References

[1] J. Keutchayan, J. Ortmann, and W. Rei. Problem-driven scenario clustering in stochastic optimization. Computational Management Science, 20(1), 2023.

[2] Y. Song and J. Luedtke. An adaptive partition-based approach for solving two-stage stochastic programs with fixed recourse. SIAM Journal on Optimization, 25(3):1344–1367, 2015.

Session 12G: Mid and Long Term Energy Management *Room: Enrico Fermi (A1.133)*

Stochastic Inflows and Seasonal Water Value at Manantali

Steeven B. Affognon, Babacar M. Ndiaye, Pierre Mendy, Cheikh M. F. Kebe, Sandrine Charousset, Omar Goudiaby, Thomas Ouillon

1Laboratory Mathematics of Decision and Numerical Analysis, Cheikh Anta Diop University of Dakar, Dakar, Senegal 2Department of Mathematics, University of Nairobi, Nairobi, Kenya, belvinos@gmail.com 3Laboratoire Eau, Energie, Environnement et procédés Industriels (LE3PI), Université Cheikh Anta Diop de Dakar (UCAD), Dakar, Senegal 4Centre de Tests des Systèmes Solaires (CT2S), Dakar, Senegal 5Électricité de France SA, Paris, France 6Hydrology Leïdi Laboratory of Gaston Berger University of Saint-Louis, Senegal

Keywords: hydropower, stochastic inflow modelling, seasonal storage valuation, stochastic dual dynamic programming, Hamilton–Jacobi–Bellman equation, Senegal river, plan4res.

Seasonal hydropower operation requires release decisions that account for uncertain future inflows and the opportunity cost of stored water. We assess the Modelling and Optimisation of Stochastic Seasonal Hydropower Storage (MOSSHOOS) framework for the Manantali reservoir on the Senegal River, using 660 monthly energy-inflow observations from 1961–2015. The framework couples a discrete and a continuous representation of the same marginal economic object. On the discrete side, a coarse Haar/Fourier seasonal decomposition and blocked residual resampling generate candidate inflow paths, which are reduced to ten weighted annual scenarios and fed into a stochastic dual dynamic programming (SDDP) reservoir model; we show that the active SDDP cut slope is exactly the shadow price of an additional unit of stored energy. On the continuous side, a seasonal square-root (Cox– Ingersoll–Ross) diffusion for inflow drives a Hamilton–Jacobi–Bellman formulation whose storage derivative $-Vs$ is the corresponding continuous-time water value; we derive the associated marginal release threshold and a Haar storage-basis approximation-rate result. As an independent economic benchmark, we repair a solver-interface defect in the open-source plan4res Seasonal Storage Valuation chain (an integer overload of set_par was silently discarding the LP-relaxation flag needed for Benders duals) and recompute its cut slopes with HiGHS. Scenario reduction improves the energy score by 47.6% relative to independent monthly sampling, but the marginal distribution is rejected by a Kolmogorov–Smirnov test and the lower-tail (p10) relative error remains 103.9%, indicating inadequate dry-tail calibration; the time-varying Feller ratio of the CIR model also falls below one in the driest months, so strict positivity cannot be assumed by the HJB solver. Consistent with this tail limitation, 50 paired Monte Carlo evaluations show no statistically detectable cost or deficit advantage of the stochastic policy over its deterministic-equivalent counterpart at the archived operating configuration (95% bootstrap intervals for both paired cost and deficit differences contain zero). The repaired plan4res benchmark yields a dominant marginal water value of 81.66 USD/MWh, within 3% of the independently computed MOSSHOOS mean of 79.25 USD/MWh, supporting economic consistency of the marginal object across the two control representations without assuming a false alignment between their respective 18 and 12 stages. These results support a coherent stochastic inflow-to-water-value methodology, cross-validated by two independent control formulations and an independent open-source economic benchmark, while showing that hydrological tail representation—not economic scaling or algorithmic design—is the principal bottleneck to address before operational deployment. We outline the extensions required for operational use: deeper multiscale scenario construction informed by a companion daily wavelet analysis, a state-constrained jump-diffusion treatment of persistent drought, and coupling to the wider OMVS cascade and to joint inflow–solar–wind–load scenarios relevant to the Senegalese power system.

References

[1] M. V. F. Pereira and L. M. V. G. Pinto, “Multi-stage stochastic optimization applied to energy planning,” Math. Program., vol. 52, pp. 359–375, 1991. [2] J. C. Cox, J. E. Ingersoll, and S. A. Ross, “A theory of the term structure of interest rates,” Econometrica, vol. 53, no. 2, pp. 385–407, 1985. [3] D. Beulertz, S. Charousset, D. Most, S. Giannelos, and I. Yueksel-Erguen, “Development of a modular framework for future energy system analysis,” in Proc. 54th Int. Universities Power Engineering Conf. (UPEC), 2019, pp. 1–6. [4] S. B. Affognon, B. M. Ndiaye, P. Mendy, and C. M. F. Kebe, “From daily fluctuations to annual hydrological cycles: a wavelet-based analysis of nonstationary seasonality in Senegal River hydropower inflows,” arXiv:2608.23470 [stat.AP], 2026. [5] S. B. Affognon, B. M. Ndiaye, P. Mendy, and C. M. F. Kebe, “InFlow: entropic stochastic dual dynamic programming with HJB cross-certification for seasonal energy storage,” arXiv:2609.11339 [math.OC], 2026.

Session 8B: Invited session: Linearly Controlled Systems *Room: Henri Poincaré (Amphi I)*

Approximate Controllability on the Group of Volume-preserving Diffeomorphisms

Andrei Agrachev[1], Bettina Kazandjian[2]

[1]SISSA, Italy, agrachev@sissa.it
[2]Sorbonne Universite, France, bettina.kazandjian@sorbonne-universite.fr

Keywords: Volume-preserving diffeomorphisms, Controllability, Ensemble of points

Let $\mathbb{T}^d = \mathbb{R}^d/(2\pi\mathbb{Z}^d)$ be the d-dimensional torus. Translations of the torus are generated by constant vector fields. Let f be a non-constant divergence-free vector field on $\mathbb{T}^d$. It generates a one-parameter group of volume-preserving diffeomorphisms. In this paper, we try to understand which transformations of $\mathbb{T}^d$ can be reached if we perturb f by constant fields (with the constant depending on time). Indeed, we study the flows generated by the control system $\dot{x} = f(x) + u(t)$, and we characterize the subgroup of volume-preserving diffeomorphisms that can be approximated by such flows when the control parameter varies within the class of piecewise constant functions. The most accurate results concern d = 2 and d = 3. Then we study the consequences of these results on the control of ensembles of points evolving under the same controlled dynamics.

References

[1] A. Agrachev, B. Kazandjian, Approximate controllability on the group of volume-preserving diffeomorphisms, preprint, arXiv, 2026.

Session 12C: Invited session: Geometric Control and Optimization *Room: André Marie Ampère (Amphi II)*

Second-order optimality conditions in geometric optimal control

Andrei A. Agrachev, Michele Motta2.

1SISSA, Trieste, Italy agrachevaa@gmail.com, mimotta@sissa.it

Invited session: Geometric Control and Optimization

Keywords: geometric optimal control, second-order optimality conditions, strong Legendre condition, conjugate points, nonlinear control systems

In the framework of optimal control, the Pontryagin Maximum Principle (PMP) is a well established and universal tool used to solve optimal control problems. However, PMP gives only necessary conditions of optimality and a trajectory satisfying them need not to be optimal in general. Moreover, even if some arc of such a trajectory is known to be optimal, PMP alone does not determine where the trajectory ceases to be an optimizer. For these reasons, secondorder optimality conditions were introduced. I this talk, I will discuss some recent optimality conditions of second order that we obtained [1], particularly for piecewise regular extremal arcs, that is concatenation of extremal arcs satisfying the strong Legendre condition. Finally, I will discuss some simple examples to illustrate how our optimality conditions work in practice.

References

[1] Andrei A. Agrachev and Michele Motta. Second order optimality conditions for piecewise regular extremals in Optimal Control. 2026. arXiv: 2607.10434 [math.OC]. url: https: //arxiv.org/abs/2607.10434.

Session 8J: Optimal Tariffs *Room: Emmy Noether (A1.140)*

Combining Optimal LEZ Scheduling and EV Purchase Incentives for Transport Policy Backcasting

Asmae Alami[1], Vinith Lakshmanan[1], Antonio Sciarretta[1]

[1]IFP Energies Nouvelles, France, asmae.alami-ouazzani@ifpen.fr
vinith-kumar.lakshmanan@ifpen.fr, antonio.sciarretta@ifpen.fr

Keywords: optimal control, policy backcasting, multi-objective optimization, NSGA-II

Reaching greenhouse gas reduction targets and improving air quality in the transport sector requires combining several policy levers, whose joint impact is far from straightforward. Low emission zones (LEZ), which restrict access to the most polluting vehicles, and electric vehicle (EV) purchase incentives are two policies that have been extensively studied in the literature. Yet, they are traditionally evaluated separately, ignoring their simultaneous interactions and effects on the vehicle fleet. This work proposes a backcasting methodology to support the combined implementation of these two policies. Rather than forecasting the effects of a given policy, backcasting starts from a set of predefined targets (e.g. environmental, budgetary) and searches for the combination of policy levers that achieves them. In our work, we use the Îlede-France region as a case study, where the evolution of the vehicle fleet is modeled through its renewal dynamics, and the LEZ implementation schedule together with the EV purchase incentive are formulated as discrete time-dependent control variables. The resulting optimal control problem is large scale and nonlinear, which adds even more complexity. Three criteria are considered in this study: (i) an emission indicator (CO2, NOx or PM), (ii) the budget which is captured by the total cost of EV purchase incentives, and (iii) a societal indicator given by the number of vehicles whose replacement is forced by the LEZ schedule, reflecting the acceptability of LEZ restrictions. Any of the three criteria can be minimized as an objective, or set as a constraint whenever maximum thresholds can be chosen by decision-makers. The criteria set as an objective can be optimized jointly in a multi-objective setting, or combined into a single mono-objective criterion, with weights chosen by the decision-maker according to their priorities. The results presented in this study focus on the unconstrained multi-objective formulation, solved with the genetic algorithm NSGA-II, as a first step to map out the Pareto front of achievable trade-offs. Once the decision-maker has situated their priorities on this front, thresholds or weights can be fixed for a more targeted formulation. The purpose of the Île-de-France case study is to illustrate this Backcasting methodology as a decision-support tool, rather than to recommend a specific policy setting. It shows how these trade-offs can guide the combination of several transport policy instruments, under multiple requirements. While NSGA-II provides a first set of results, further work is needed to explore other resolution methods for this complex optimal control problem, and to improve computational efficiency.

Session 15F: Combinatorial Optimization & Applications *Room: Ada Lovelace (A1.128)*

Comparing Linear Programming and Mean-Payoff Games via Oriented Matroids

Xavier Allamigeon[1], Yiyuan Chen[1], Stéphane Gaubert[1], Xavier Goaoc[2]

[1]CMAP, École Polytechnique & INRIA Saclay, France
firstname.lastname@inria.fr
[2]LORIA, Université de Lorraine & INRIA Nancy, France, xavier.goaoc@loria.fr

Keywords: Linear programming, mean-payoff games, oriented matroids, tropical geometry, parity games.

Linear programming, mean-payoff games, and parity games raise long-standing questions in algorithmic complexity. Tropical geometry connects these problems: mean-payoff games are equivalent to feasibility problems for tropical linear inequalities [1]. In particular, a strongly polynomial combinatorial simplex algorithm for linear programming would yield a strongly polynomial algorithm for mean-payoff games [2]. These connections lead us to study the underlying combinatorial structures using oriented matroids, which encode the determinant signs on which combinatorial pivoting rules depend. In this talk, we compare the families of oriented matroids arising from classical, tropical, and parity point configurations in general position. For every fixed rank r $\geq$2, we prove that the tropical and parity families of realizable oriented matroids on n labeled elements have respective sizes

$$n^{(r-1)^2 n+o(n)} \qquad \text{and} \qquad n^{(r-1)n+o(n)} \quad \text{as } n \to \infty.$$

Although tropically realizable oriented matroids form a subclass of classically one, both families have size $n^{(r-1)^2 n+o(n)}$, whereas parity realizable oriented matroids form a family of size $n^{(r-1)n+o(n)}$ that grows strictly more slowly when r $\geq$3. These results give a quantitative comparison of the combinatorial expressivity of linear programming, mean-payoff games, and parity games. Beyond the oriented matroid results, we show that Santos's proposed estimate $(nr)^{\Theta(nr)}$ for the number of regular subdivisions of $\Delta_{n-1} \times \Delta_{r-1}$ cannot hold uniformly in n and r [3].

References

[1] M. Akian, S. Gaubert, and A. Guterman. "Tropical Polyhedra Are Equivalent to Mean Payoff Games". In: International Journal of Algebra and Computation 22.1, 1250001 (2012).

[2] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. "Combinatorial Simplex Algorithms Can Solve Mean Payoff Games". In: SIAM Journal on Optimization 24.4 (2014).

[3] F. Santos. "The Cayley Trick and Triangulations of Products of Simplices". In: Integer Points in Polyhedra—Geometry, Number Theory, Algebra, Optimization. Vol. 374. Contemporary Mathematics. American Mathematical Society, 2005.

Session 5D: Invited session: Black-box Optimization *Room: Lise Meitner (A1.116)*

Optimization-Oriented Surrogate Frameworks for Wind Farm Layout Optimization

Ekaterina ANDREEVA, Claudia D'Ambrosio, Ahlam MAGNOUCHE

1 LIX, CNRS, Ecole polytechnique, Institut Polytechnique de Paris, Palaiseau, France
ekaterina.andreeva@lix.polytechnique.fr, dambrosio@lix.polytechnique.fr 2 TotalEnergies OneTech, Palaiseau, France
ekaterina.andreeva@totelenergies.com, ahlam.magnouche@totalenergies.com

Keywords: Surrogate-based optimization, black-box optimization, machine learning, mixedinteger optimization, wind farm layout optimization

Many real-world optimization problems involve objective functions whose evaluation is computationally expensive, making direct optimization impractical. Surrogate-based optimization addresses this challenge by replacing costly evaluations with predictive models learned from a limited number of samples. Although surrogate models are commonly assessed through predictive accuracy metrics, their suitability for optimization can be overlooked. This work investigates optimization-oriented surrogate-assisted frameworks for expensive black-box optimization, with a focus on Wind Farm Layout Optimization (WFLO). The objective is not to develop new surrogate models, but rather to study how surrogate representations, optimization formulations, and sequential sampling strategies can be combined to efficiently guide the search towards high-quality solutions. Particular attention is paid to the interaction between model accuracy, and optimization tractability. The proposed research explores different surrogate families and optimization schemes within iterative design-of-experiments frameworks inspired by Efficient Global Optimization and Sequential Model-Based Optimization [1]. The long-term goal is to identify surrogate-optimization combinations capable of scaling to large and expensive optimization problems while maintaining strong optimization performance.

References

[1] D.R. Jones, M. Schonlau, W.J. Welch, Efficient Global Optimization of Expensive BlackBox Functions, Journal of Global Optimization, 13, 455-492 (1998), DOI:1008306431147.

Session 8J: Optimal Tariffs *Room: Emmy Noether (A1.140)*

Towards Dynamic and Adaptive Tariffs in Electricity Distribution Networks

Panagiotis Andrianesis[1], Vasileios Cheimonidis[1], Alexandre Manso de Zuniga[1]

[1]Mines Paris - PSL University, France, panagiotis.andrianesis@minesparis.psl.eu

Keywords: dynamic/adaptive distribution network tariffs, optimization and learning, spatiotemporal marginal costs, network sensitivities, optimal grid - distributed energy resource coordination.

The estimation of Distribution Locational Marginal Costs (DLMCs), i.e., the marginal costs for serving real and reactive power demand at a certain time and location, can serve as a means for efficient coordination of distributed energy resources in active distribution networks [1, 2]. DLMC computation can benefit from the computation of voltage and current sensitivities with respect to nodal real and reactive net demand, yet detailed network models are often unavailable in practice. Inspired by [3], we derive a parameterized formulation of these sensitivities that explicitly captures their dependence on net demand, based on the LinDistFlow model. Building on this formulation, we propose an optimization based approach that learns network parameters and bias terms from measurement data or high-fidelity power flow outputs. Numerical results on standard test feeders, and on an actual distribution feeder demonstrate the effectiveness of our data driven approach and its potential to support dynamic and adaptive distribution network tariffs [4]. Additional results on an actual distribution feeder that employ an iterative coordinate descent algorithm illustrate that the coordinated charging of electric vehicles, which adapt their charging profiles to these tariffs, can achieve most of the system-wide efficiency gains in just a few iterations.

References

[1] P. Andrianesis, M. Caramanis and N. Li, Optimal Distributed Energy Resource Coordination: A Decomposition Method Based on Distribution Locational Marginal Costs," IEEE Transactions on Smart Grid, 13(2), 1200–1212, 2022.

[2] P. Andrianesis, M. Caramanis, Distribution Network Marginal Costs: Enhanced AC OPF Including Transformer Degradation, IEEE Transactions on Smart Grid, 11(5), 3910–3920, 2020.

[3] B. Taheri, R. K. Gupta, and D. K. Molzahn, Optimizing parameters of the LinDistFlow power flow approximation for distribution systems, IEEE Transactions on Smart Grid, 17(2), 1029–1042, 2026.

[4] V. Cheimonidis, P. Andrianesis, Learning voltage and current sensitivities for marginal cost estimation in active distribution networks, Proc. 22nd International Conference on the European Energy Market (EEM26), Trondheim, Norway, 22 - 24 June (2026).

This work has benefited from the support of the European Research Council (ERC) Starting Grant CRISP under Agreement 101165616, and from insightful discussions with Michael Caramanis (Boston University).

Session 12A: Invited session: Applications of optimization to power systems and networks *Room: Marie Curie (Auditorium)*

Single-Level Reformulations for Network Pricing under Wardrop Equilibrium with Separable Affine Costs

Miguel F. Anjos, José Neto

1 Group for Research in Decision Analysis (GERAD), Montréal, Canada 2 School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, United Kingdom, anjos@stanfordalumni.org 3 SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France, jose.neto@telecom-sudparis.eu

Keywords: bilevel optimization, network pricing, Wardrop equilibrium, single-level reformulation, column-and-constraint generation

We consider the problem of setting tolls on a subset of the arcs of a congested network so as to maximize the revenue of a network authority, anticipating that users route according to a Wardrop equilibrium. This problem is naturally cast as a bilevel program in which the leader sets tolls while the followers respond by choosing an equilibrium routing. We focus on the important special case of separable affine arc cost functions. Under this assumption, Wardrop equilibria admit Beckmann's convex characterization, from which we derive three new single-level reformulations of the bilevel problem, exploiting both Beckmann's formulation and its dual: an arc-based formulation, which is exact, and two path-based formulations, solved by column-and-constraint generation. Beyond the reformulations, we further derive from Beckmann's potential an upper bound on the optimal revenue that can be used to certify optimality gaps across all formulations. An extensive computational study on instances derived from the Sioux Falls and Eastern Massachusetts networks compares our formulations with the recent mixed-integer nonlinear formulation of Beck, Labbé and Schmidt [1]. The arc-based formulation solves substantially more instances to proven optimality (within a time limit of 1800s) than the formulation of [1] and, unlike it, returns a feasible solution on every instance, remaining effective on the largest 74node network, on which [1] solves few instances, if any. The magnitude of the upper bounds on the tolls proves to be a major driver of tractability: the genuinely hard core of the benchmark emerges once the regulatory cap on tolls is removed. The two path-based formulations, used as heuristics, return a solution within 1% of the optimum on every instance whose optimal value is known, often in seconds where exact formulations reach the time limit.

References

[1] Y. Beck, M. Labbé, M. Schmidt, A toll-setting problem with robust Wardrop equilibrium conditions under budgeted uncertainty, Optimization Online, 2024, optimizationonline.org/?p=26949.

Session 12K: Multiobjective Optimization and Evolutionary Computation Part I *Room: Loggia (A1 - first floor)*

Online Algorithm Selection: from Foundations to Applications

Denis Antipov[1], Carola Doerr[1], Diederick Vermetten[1]

[1]Sorbonne Université, CNRS, LIP6, France,
denis.antipov@lip6.fr, carola.doerr@lip6.fr, diederick.vermetten@lip6.fr

Keywords: Random Search Heuristics, Algorithm Selection, Runtime Analysis, Evolutionary Algorithms, Theory

The effectiveness of optimization algorithms depends drastically on the problem at hand, and often it can change during the optimization process. This implies that choosing different algorithms for different stages of search might outperform any single algorithm that is used from the beginning until the end. However, finding a good strategy of switching between algorithms from a given portfolio for a specific problem is a non-trivial task. In this work, we provide the first foundations for online algorithm selection (OAS) and show an example where its superiority over single-algorithm approaches is well-pronounced. Namely, we consider a setting when we have a portfolio of two algorithms, the (1 + 1) EA and the (1 + (λ, λ)) GA, and we solve OneMax benchmark by starting with the (1 + 1) EA and switching to the (1 + (λ, λ)) GA when close to the optimum. While both algorithms from the portfolio have $\Omega(n/\sqrt{\log n})$ expected runtime (or by sub-polylogarithmic factors better), the OAS can find the optimum in O(n log log n) expected time, which is asymptotically faster. We then provide a broad empirical study which confirms that our observation that these two algorithms demonstrate different effectiveness at different distances from the optimum. It also indicates that the proposed switching mechanism is not specific to the OneMax benchmark, but rather general. We also present the current work-in-progress, where we explore more complicated mechanisms of switching between algorithms such as reinforcement learning and bandit-based approaches. We show that these approaches can be used to learn the best switching strategy for a given problem, and we provide preliminary results on their performance.

Session 15G: Data Science, Learning and the intersection with Optimization *Room: Enrico Fermi (A1.133)*

Deterrence of Manipulation through Preference Elicitation in Personalized Incentive Policies

Andrea Araldo, Alessa Mayer, Ravi Seshadri, Amir Brudner

1Institut Polytechnique de Paris, France, araldo@telecom-sudparis.eu 2Technical University of Denmark, Denmark, ravse@dtu.dk 3Massachusetts Institute of Technology, United States, amirb@mit.edu

Keywords: Games, Manipulation, Incentive Policies, Preference Elicitation, Learning

Demand-management policies aim to shift behavior toward sustainable choices, such as goods or transport modes. Personalized incentive policies tailor subsidies to agents' needs, but they also enable strategic manipulation: an agent may imitate a less responsible profile to obtain larger incentives. When this is profitable remains poorly understood. We model the repeated interaction as a Stackelberg game. Prior work often represents manipulation as strategic misreporting: a follower supplies fake payoffs or best responses, and the leader adapts its commitment accordingly [1]. We instead place preference elicitation inside the game: the regulator infers a profile from a history of random-utility choices and applies its optimized incentive. This separates the agent's true, imitated, and regulator-inferred profiles. To imitate another profile, the agent must repeatedly choose as that profile would rather than follow its true preferences, thereby incurring a Cost of Manipulation (CoM). We decompose the advantage from manipulation into a Gain from Manipulation (CoM) (obtained from shifting the distribution of future incentives towards higher values) minus the CoM. Information-theoretic arguments upper-bound the GoD in terms of the regulator's ability to distinguish profiles from past choices. Comparing this bound with the CoD yields a threshold on the interaction length: below this threshold, manipulation is guaranteed to be unprofitable. We derive a counter-intuitive analytical result: when preference elicitation is less effective—i.e., when the regulator has greater difficulty distinguishing agent profiles from their past behavior—manipulation is more strongly deterred. A Monte Carlo case study calibrated on London travel choices [2] shows that when the agent can simply report their profile to the regulator, the agent is always better off manipulating. However, in the most realistic case in which the regulator infers the agent's profile via preference elicitation, the manipulation would entail a too high Cost of Manipulation, and the agent is better off being honest. The framework can thus stress-test personalized incentives before deployment.

References

[1] G. Birmpas et al., "Optimally deceiving a learning leader in Stackelberg games," NeurIPS, 33:20624–20635, 2020.

[2] T. Hillel, M. Z. E. B. Elshafie, and Y. Jin, "Recreating passenger mode choice-sets for transport simulation: A case study of London, UK," Proc. ICE–Smart Infrastructure and Construction, 171(1):29–42, 2018.

Session 12A: Invited session: Applications of optimization to power systems and networks *Room: Marie Curie (Auditorium)*

Benders decomposition with cut selection for long-term generation and storage capacity expansion problems

Ayse N. Arslan, Olivier Baude, Boris Detienne, Aurélien Froger, Rodolphe Griset, Cécile Rottner, Marius Roland, Frédéric Semet, Mariam Sangaré

1Inria, Centre Inria de l'Université de Bordeaux, Talence, France 2Université de Bordeaux, Institut de Mathématiques de Bordeaux, Talence, France 3EDF R&D, OSIRIS, Palaiseau, France 4Inria, Centre Inria de l'Université de Lille, Villeneuve-d'Ascq, France 5Centrale Lille, Villeneuve-d'Ascq, France

Keywords: Capacity expansion, Benders decomposition, cut selection, energy systems

The expansion planning (EP) problem in energy systems optimization aims to determine long-term investment decisions while ensuring the feasibility of short-term system operations. Long-term decisions concern investments in generation and storage capacities, as well as capacity decommissioning and lifetime extension, whereas short-term operational decisions involve generation and storage scheduling subject to supply–demand balance, technical, and network constraints [1]. We consider Benders decomposition to exploit this natural separation between investment and operational decisions. However, accurately representing short-term operations over a long planning horizon leads to large operational problems, particularly in the presence of timecoupling constraints. To address this issue, we rely on a temporal decomposition that separates the operational problem into smaller subproblems [2]. At each Benders iteration, these subproblems can generate several candidate cuts for the current master-problem solution, raising the question of which cuts should be added to the master problem. In this talk, we present this temporal decomposition approach and then focus on cut selection, investigating how different selection policies affect the performance of the Benders decomposition algorithm. We review several strategies from the literature [3] and discuss theoretical properties of Benders cuts. Finally, we provide a numerical analysis on capacity expansion instances derived from the ERAA dataset to compare the performance of these selection strategies.

References

[1] C. L. Lara, D. S. Mallapragada, D. J. Papageorgiou, A. Venkatesh, I. E. Grossmann, Deterministic electric power infrastructure planning: Mixed-integer programming model and nested decomposition algorithm, European Journal of Operational Research, 271(3), 2018.

[2] A. Jacobson, F. Pecci, N. Sepulveda, Q. Xu, J. Jenkins, A computationally efficient Benders decomposition for energy systems planning problems with detailed operations and timecoupling constraints, INFORMS Journal on Optimization, 6(1):32–45, 2024.

[3] M. Hosseini, J. Turner, Deepest cuts for Benders decomposition, arXiv:2110.08448, 2024.

Session 15F: Combinatorial Optimization & Applications *Room: Ada Lovelace (A1.128)*

Through the G-Polymatroid Glass and What Unit Commitment Finds There

Hélène Arvis[1,2,3], Olivier Beaude[1], Nicolas Gast[2], Stéphane Gaubert[3], Bruno Gaujal[2]

[1]EDF Lab Paris-Saclay, OSIRIS department, France helene.arvis@edf.fr
[2]Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG
[3]Inria, and CMAP, École polytechnique, IP Paris, CNRS

Keywords: Combinatorial Optimization, Aggregation, Generalized Polymatroids, Unit Commitment Problem

When considering substantial Unit Commitment (UC) problems, a common practice for reducing the problem size is that of aggregating similar units in order to reduce the computational complexity of the problem [1]. However, naive aggregation methods generally result in infeasibilities on the individual level. Nonetheless, it has been shown that Electric Vehicle (EV) flexibility sets are in fact generalized polymatroids [2, 3], or g-polymatroids, a class of polytopes introduced by Frank [4] and further studied in the context of discrete convexity by Danilov and Koshevoy [5]. As gpolymatroids are stable by Minkowski sum and allow optimization of linear objectives via the greedy algorithm, this observation ensures an exact aggregation of EVs and an efficient resolution. We explore the other components of the UC problem as reported by the European Resource Adequacy Assessment (ERAA) and show that other storage and flexibility means, such as batteries or heat networks, also have g-polymatroid constraints. We also evoke more difficult models, such as detailed hydropower, which do not exhibit the g-polymatroid structure (at least not in their most natural formulation).

References

[1] Jelle Meus, Kris Poncelet, and Erik Delarue. Applicability of a Clustered Unit Commitment Model in Power System Modeling . In: IEEE Transactions on Power Systems 33.2 (2018), pp. 2195 2204. doi: 10. 1109/TPWRS.2017.2736441.

[2] Karan Mukhi, Georg Loho, and Alessandro Abate. Exact Characterization of Aggregate Flexibility via Generalized Polymatroids . In: IEEE Transactions on Smart Grid 17.3 (2026), pp. 2131 2142. doi: 10. 1109/TSG.2026.3659513.

[3] Hélène Arvis et al. Integrating Aggregated Electric Vehicle Flexibilities in Unit Commitment Models using Submodular Optimization . In: 2026 European Control Conference (ECC). ISSN: 2996-8895. 2026, pp. 760 766.

[4] A. Frank. Generalized Polymatroids . In: Finite and Infinite Sets. Elsevier, 1984, pp. 285 294. doi: 10. 1016/B978-0-444-86893-0.50021-8.

[5] Vladimir I Danilov and Gleb A Koshevoy. Discrete convexity and unimodularity I . In: Advances in Mathematics 189.2 (2004), pp. 301 324. doi: 10.1016/j.aim.2003.11.010.

Session 15C: Energy Systems *Room: André Marie Ampère (Amphi II)*

Nodal-Aware Flow-Based Capacity Calculation Through Bilevel Inverse Optimization

Pierre Arvy, Florin Capitanescu, Amélie Lambert, Daniel Porumbel, Nicolas Omont

1Artelys, Paris, France, pierre.arvy@artelys.com 2Conservatoire national des arts et métiers, Paris, France, 3Luxembourg Institute of Science and Technology, Esch-sur-Alzette, Luxembourg

Keywords: Bilevel optimization, inverse optimization, flow-based market coupling, electricity markets, decomposition

Flow-based market coupling clears the European day-ahead electricity market subject to a zonal approximation of transmission constraints constructed ex ante. Because this approximation is imperfect and zonal net positions do not uniquely determine nodal injections, the resulting dispatch may violate physical limits and require costly ex post remedial actions during congestion-management stage. We ask whether nodal security can instead be incorporated into flow-based capacity calculation while leaving the standard zonal clearing problem and its pricing rule unchanged. We formulate capacity calculation stage as a bilevel inverse-optimization problem. The upper level chooses admissible flow-based parameters, which define the zonal approximation of transmission constraints, and a nodally secure dispatch. The lower level represents conventional zonal market coupling. Since both levels maximize the same welfare function, the formulation seeks the best secure outcome that an admissible flow-based feasibility domain can induce. Our decomposition exploits a structural partition of the lower-level Karush–Kuhn–Tucker (KKT) conditions. Resource-side conditions link dispatch quantities to zonal prices independently of the flow-based parameters and strengthen a hybrid zonal–nodal master problem. The master thus selects a secure target consistent with zonal merit order. An inverse-optimization subproblem then calibrates the flow-based parameters, with regularization toward conventional capacitycalculation estimates, and reconstructs the omitted network-side conditions. Compared with a conventional decomposition using a purely nodal master, this construction excludes targets that no zonal price vector can support, removing a major source of inverse-subproblem infeasibility. Feasibility of the inverse subproblem completes the KKT system and certifies that the target is globally optimal for the unchanged zonal clearing problem under the calibrated parameters. Global solution of the master additionally certifies global optimality for the bilevel problem. We characterize supporting zonal prices, give a rank condition for uniqueness, and identify structural causes of reconstruction infeasibility. Experiments on a stylized network and the IEEE 118-bus system show that calibrated domains can eliminate physical violations at market clearing and reduce total operating cost when ex post congestion-management costs are included. The framework thus provides both a capacity-calculation method and explicit certificates that nodally secure outcomes can be implemented through zonal market coupling.

Session 8A: Learning Dynamics and Optimization *Room: Marie Curie (Auditorium)*

Survival of Dominated Strategies under Standard No-Regret Dynamics

Galit Ashkenazi-Golan, Domenico Mergoni Cecchelli, Panayotis Mertikopoulos, Edward Plumb, Bassel Tarbush, Yannick Viossat

1London School of Economics, UK, g.ashkenazi-golan@lse.ac.uk e.plumb@lse.ac.uk 2CUNEF Universidad, Spain, domenico.mergoni@cunef.edu 3CNRS, France panayotis.mertikopoulos@imag.fr 4Oxford University, UK bassel.tarbush@economics.ox.ac.uk 5Université Paris Dauphine-PSL, France viossat@ceremade.dauphine.fr

Keywords: Game theory, learning, no-regret, dominated strategy, survival

Many no-regret dynamics are known to eliminate strictly dominated strategies. Nevertheless, we show on a 3×2 example that a standard no-regret dynamics, the projected gradient algorithm, may fail to eliminate pure strategies strictly dominated by other pure strategies. This holds to an open set of games, and extends to the optimistic variant of the projected gradient algorithm. This may occur because, by contrast with other major no-regret dynamics such as the exponential weight algorithm or potential based no-regret dynamics à la Hart and Mas-Collel, the projected gradient algorithm is not a perturbed version of continuous fictitious play. Our results show that, though achieving no-regret seems harder than avoiding irrational behaviour, examples of no-regret distributions putting weight on dominated strategies are not mere intellectual curiosities: standard no-regret algorithms may really lead to such distributions of play.

Session 15B: Learning & Neural Networks *Room: Henri Poincaré (Amphi I)*

Benchmarking Deep Learning Architectures for Industrial Non-Intrusive Load Monitoring: Performance, Training Efficiency, and Granularity

Laiba Asmatullah, Valentin Colliard, Alizée Mesnard

1ShiftElec, France, valentin.colliard@shiftelec.com 2SAMOVAR, Telecom SudParis, Institut Polytechnique de Paris, 91120 Palaiseau, France, laiba.asmatullah@telecom-sudparis.eu, alizee.mesnard@telecom-sudparis.eu

Keywords: Non-Intrusive Load Monitoring, Energy Disaggregation, Deep Learning, Data Granularity, Energy Management

In the context of global decarbonization, optimizing industrial electricity consumption is a critical lever for reducing greenhouse gas emissions. Non-Intrusive Load Monitoring (NILM) estimates individual load consumption from an aggregate meter, offering a cost-effective means to gain this visibility. Although deep learning has advanced residential NILM, industrial deployment remains difficult because machine signatures overlap, operating cycles are long and heterogeneous, variable-speed loads are common, and operating states are strongly imbalanced. This difficulty is compounded by the scarcity of publicly available industrial datasets, which in turn limits the availability of standardized benchmarks for fairly comparing methods across machine types. We propose an industrial benchmark of representative deep learning methods for NILM, including convolutional, recurrent, multi-task, attention-based, Transformer, and generative architectures. We compare seventeen implementations on five industrial machines from the HIPE dataset [1], spanning binary, multi-state, and continuously variable behavior. We organize our analysis around three axes: performance (measured via Total Energy Correctly Assigned, TECA), learning speed, and sampling resolution. Results show strong machine dependence: BiLSTM gives the highest mean TECA overall (0.543), while BiGRU leads on Chip Press (0.735) and Attention CNN on the High Temperature Oven (0.726). A separate runtime experiment shows that a sequence-to-sequence CNN trains roughly 13x faster than BiGRU and 21x faster than Energformer under our measured setting. A sampling-resolution sweep from 5 s to 1 h across ten machines shows 5 s gives the highest mean TECA (0.570), though the optimum remains machine-specific. These results support architecture and sensing choices conditioned on machine dynamics rather than a universal model or sampling interval.

References

[1] Simon Bischof, Holger Trittenbach, Michael Vollmer, Dominik Werle, Thomas Blank, and Klemens Böhm. 2018. HIPE: An Energy-Status-Data Set from Industrial Production. In Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy '18). Association for Computing Machinery, New York, NY, USA, 599–603. https://doi.org/10.1145/3208903.3210278

Session 15B: Learning & Neural Networks *Room: Henri Poincaré (Amphi I)*

Using statistical learning in Operations Research: Generalization Bounds and Diffusion Models

Pierre-Cyril Aubin-Frankowski

Keywords: machine learning, combinatorial optimization, statistical learning, diffusion models

Many real-world decision problems require solving, again and again, combinatorial optimization instances drawn from a common distribution. Instead of solving each instance independently, a recent line (COAML) of structured learning methods, with many empirical achievements, exploits the statistical regularity by learning policies that pair a statistical model with a tractable combinatorial oracle. We here analyze smoothed policies: adding controlled random perturbations to the direction used by the linear oracle yields a differentiable surrogate risk and improves generalization. Our main contribution is a generalization bound that decomposes the excess risk into i) perturbation bias, ii) statistical estimation error, and iii) optimization error. The perturbation bias is controlled by the fan-crossing probability, a new geometric quantity measuring the likelihood that a perturbation changes the oracle solution. The statistical model is then eventually replaced by a denoising diffusion model, allowing to use probabilities over continuous sets despite the discrete setting of combinatorial problems.

Based on ongoing work with Yohann De Castro (ICJ Lyon) and Axel Parmentier (CERMICS, ENPC) and on https://arxiv.org/abs/2407.17200 (to be accepted at SIOPT)

Session 8F: Solving games and Equilibrium problems *Room: Ada Lovelace (A1.128)*

Computing Projected solutions of Nash games : from theory to electricity markets

Didier Aussel, Jauny, Asrifa Sultana, Shivani Valecha

1Lab. Promes UPR CNRS 8521, University of Perpignan, France, aussel@univ-perp.fr 2University of Rising Stars, USA, second.author@rising-stars.edu, third.author@rising-stars.edu

Keywords: Nash games, projected solutions, electricity market

The concept of projected solution for nonself generalized Nash games and quasi-variational inequalities has been defined in [1] and this approach has attracted attention in different context (see e.g. [2, 3, 4, 5, 6, 7, 12]). Nevertheless the computation of projected solutions has been only recently investigated ([8, 9, 10, 11]). Our aim in this talk is to present a new numerical approach and its application for the modelling of bid processes in electricity markets.

References

[1] D. Aussel, A. Sultana, V. Vetrivel, On the existence of projected solutions of quasivariational inequalities and generalized Nash equilibrium problems, J. Optim. Theory Appl. 170 (2016), no. 3, 818-837.

[2] I. Benedetti, S. Lorenzini, Projected solutions for generalized quasivariational inequalities without compact constraint multimap, Commun. Nonlinear Sci. Numer. Simul. 129 (2024), Paper No. 107687, 19 pp.

[3] O. Bueno, J. Cotrina, Existence of projected solutions for generalized Nash equilibrium problems, Optim. Theory Appl. 191 (2021), no. 1, 344-362.

[4] M. Castellani, M. Giuli, S. Latini, Projected solutions for finite-dimensional quasiequilibrium problems, Comput. Manag. Sci. 20 (2023), no. 1, Paper No. 9, 14 pp.

[5] M. Castellani, M. Giuli, M. Milasi, D. Scopelliti, Projected solutions of generalized quasivariational problems in Banach spaces, Nonlinear Anal. Real World Appl. 76 (2024), Paper No. 104021, 12 pp.

[6] F. Guan, J. Cen, B. Chen, J.C. Yao, Existence of projected solutions for quasi-variational hemivariational inequality, Demonstr. Math. 57 (2024), no. 1, Paper No. 20230139, 7 pp.

[7] J. Cotrina, C.Calderon, Remarks on projected solutions in generalized Nash games, Oper. Res. Lett. 62 (2025), Paper No. 107323, 5 pp.

[8] G. Bigi, M. Castellani, S. Latini, An extragradient algorithm for a lifted reformulation of projected solutions for quasiequilibria, J. Global Optim. 92 (2025), no. 4, 1071-1090.

[9] G. Bigi, M. Castellani, S. Latini, Nonsmooth techniques for computing projected solutions of quasiequilibria via gap functions, J. Optim. Theory Appl. 208 (2026), no. 1, Paper No. 39, 21 pp.

[10] M. Ramazannejad, Douglas–Rachford splitting algorithm for projected solution of quasi variational inequality with non-self constraint map, Optim. Lett. 20 (2026), no. 3, 625–640.

[11] M. Bianchi, E. Miglierina, M. Ramazannejad, On projected solutions for quasi equilibrium problems with non-self constraint map, Optimization 75 (2026), no. 7, 1583–1607.

[12] C. Calderon, M. Castellani, J. Cotrina, M. Giuli, Reformulations for Projected Solutions of Generalized Games, J. Optim. Theory Appl. 204, 4 (2025).

Session 5E: Handling Large scale Stochastic Programs *Room: Antoine Lavoisier (A1.122)*

Stronger Decomposition Cuts for Stochastic Unit Commitment Using Interval Variables

Mathis Azema, Vincent Leclère, Wim van Ackooij
CERMICS, ENPC; ENPC; EDF R&D

Keywords: Unit Commitment, stochastic optimization, decomposition methods, interval variables

The Unit Commitment (UC) problem is a fundamental challenge for energy producers and network operators aiming to determine an optimal generation schedule for electricity production units while satisfying technical and economic constraints. A major source of complexity arises from uncertainty, particularly in forecasting energy demand and renewable generation. To address this issue, decomposition methods have been widely studied to efficiently handle a large number of scenarios.

Recently, a new formulation of the deterministic UC problem based on interval variables has been proposed to better model unit production decisions. Building on this approach, we develop stronger cuts for stochastic UC problems. Interval variables represent the complete online intervals of the units and therefore capture more information than the classical formulation. This additional information allows us to better approximate future operating costs within decomposition algorithms.

Through computational experiments on benchmark instances, we show that the proposed cuts improve convergence and reduce solution times. These results illustrate the benefit of interval variables for solving large-scale stochastic UC problems and their potential for decomposition methods in which decisions are adapted as uncertainty is revealed.

Session 15J: Operations Research for Food and Healthcare *Room: Emmy Noether (A1.140)*

Stochastic analysis of an emergency service with two levels sharing the servers

Guillaume Ballif[1], Christine Fricker[1]

[1]INRIA Paris, France, guillaume.ballif@inria.fr, christine.fricker@inria.fr

Keywords: Fluid limit, Differential inclusions, optimal policies

In this talk, a model for emergency department as a stochastic network with a two-level service is analyzed. In the model, a patient care pathway consists of a first consultation, followed by medical examinations and finally an exit consultation before discharge. The first level (the technical cubicles) contains the patients before or during the first consultation. The second level (a waiting zone) consists of patients undergoing examinations and patients waiting for or attending a second consultation. Each level has a finite capacity, so when the second level is full, patients must stay at the first level. A main feature of the model is the fact that first and second consultations are managed by the same pool of emergency physicians. Under a scaling assumption, for different service policies, the evolution of the number of patients in each step of the pathway is investigated. It is shown that depending on the parameters, there is a phase transition for the system between two different regimes: a regime limited by the number of emergency physicians and a regime limited by the second level capacity. Stochastic calculus from the evolution equations of the 2 or 3-dimension Markov process describing the system, martingale properties and differential inclusions are the main mathematical tools to obtain the limit of the scaled process. Using this limiting process, the various service policies are described and compared in order to determine, for each regime, the policy that best optimizes either the number of patients treated per unit of time or the average duration of a care pathway.

Session 5I: Invited Session:Neural Network approximations for high-dimensional PDEs Problems*Room: Sophie Germain (A1.139)*

Mean-Field Generalisation Bounds for Learning Controls in Stochastic Environments

Boris Baros*, Samuel Cohen*, Christoph Reisinger*

*Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK, {boris.baros, samuel.cohen, christoph.reisinger}@maths.ox.ac.uk

We consider a data-driven formulation of the classical discrete-time stochastic control problem. Our approach exploits the natural structure of many such problems, in which significant portions of the system are uncontrolled. Employing the dynamic programming principle and the mean-field interpretation of single-hidden layer neural networks, we formulate the control problem as a series of infinite-dimensional minimisation problems. When regularised carefully, we provide practically verifiable assumptions for non-asymptotic bounds on the generalisation error achieved by the minimisers to this problem, thus ensuring stability in overparametrised settings, for controls learned using finitely many observations. We explore connections to the traditional noisy stochastic gradient descent algorithm, and subsequently show promising numerical results for some classic control problems.

Session 12J: Handling dynamics *Room: Emmy Noether (A1.140)*

Generating electricity with a kite – optimization and bifurcation analysis

Antonin Bavoil[1], Romain Veltz[2]

[1]Université Côte d'Azur, France, antonin.bavoil@univ-cotedazur.fr
[2]Centre Inria d'Université Côte d'Azur, France, romain.veltz@inria.fr

Keywords: Airborne Wind Energy, Modelling, Shooting, Bifurcation Analysis

We study the numerical optimization and bifurcation analysis of an Airborne Wind Energy (AWE) system. Using kites to harness wind energy has been intensively studied in the last decade, see e.g. the survey by M. Diehl et al. in [1]. In the framework of the KEEP project (Kite Electrical Energy Production), funded by PGMO and CNRS and gathering researchers from ENSTA Bretagne and Université Côte d'Azur, our goal is to maximize the electrical power generated by a kite attached to a lever arm; flying the kite along a crosswind lemniscate drives the arm to generate electricity [2]. We build a point-mass mechanical model where a fictitious aerodynamic force steers the kite along a prescribed lemniscate-shaped cone. The resulting differential equation is expressed as a second-order ODE with two degrees of freedom (4-dimensional state space). The power maximization problem on the periodic regime is solved via a single-shooting formulation that jointly optimizes the design parameters, initial state, and period subject to periodic boundary conditions, bypassing slow asymptotic time-averaging. Finally, we use BifurcationKit.jl [3] to perform continuation on the resulting limit cycle with respect to wind speed, revealing the existence of a fold bifurcation at low wind speeds as well as confirming the local uniqueness of the limit cycle at normal operating conditions. This is an example of joint use of optimization and bifurcation theory for the numerical analysis of a dynamical system.

References

[1] U. Ahrens, M. Diehl, R. Schmehl (Eds.), Airborne Wind Energy, Green Energy and Technology, Springer, 2013, doi:10.1007/978-3-642-39965-7.

[2] A. Bavoil, J.-B. Caillau, L. Dell'Elce, A. Nême, Airborne wind energy: the KEEP project, European Control Conference (ECC), Paper WeA8.1, Reykjavík, Iceland, 2026. HAL: hal05366391.

[3] R. Veltz, BifurcationKit.jl, Technical Report, Inria Sophia-Antipolis, 2020. HAL: hal02902346.

Session 15E: Algorithms 2.0 : advances in computational methods *Room: Antoine Lavoisier (A1.122)*

Parameterized Complexity for Segment Routing

Cristina Bazgan, Morgan Chopin, André Nichterlein, Camille Richer

[1]Université Paris-Dauphine, PSL Research University, CNRS, LAMSADE, France; [2]Orange Research, France; [3]Technische Universität Berlin, Germany

Keywords: traffic engineering, segment routing, parameterized algorithms

Segment Routing is a recent network technology that helps optimizing network throughput by providing finer control over the routing paths. Instead of routing directly from a source to a target, each demand is assigned a list of waypoints called segment path, and between consecutive waypoints, the packets are routed according to traditional shortest path routing protocols. Segment Routing offers greater flexibility by allowing packets to deviate from the shortest paths, with minimal additional implementation costs for network administrators. The associated computational problem is Min Segment Routing: Given a network on n vertices with weighted and capacitated links, d traffic demands (vertex pairs), and a number k, the task is to find, for each demand, at most k waypoints such that routing the demands according to these segment paths minimizes congestion in the network. Congestion is usually measured by the maximum link utilization (MLU), that is, the largest ratio of load to capacity over the links of the network. The problem is known to be weakly NP-hard [3]. Approaches based on linear program- ming [1, 4, 2] or constraint programming [3] have been employed in previous works to solve it heuristically or speed-up solvers. We complement these practical approaches with a rigorous analysis of the computational complexity of Min Segment Routing. In particular, we study the problem within the theoretical framework of parameterized complexity: we investigate if special structures of real-world communication networks could be exploited algorithmically.

References

[1] Randeep Bhatia et al. “Optimized network traffic engineering using segment routing”. In: 2015 IEEE Conference on Computer Communications (INFOCOM). IEEE. 2015, pp. 657– 665.

[2] Hugo Callebaut, Jérôme De Boeck, and Bernard Fortz. “Preprocessing for segment routing optimization”. In: Networks 82.4 (2023), pp. 459–478.

[3] Renaud Hartert et al. “Solving segment routing problems with hybrid constraint program- ming techniques”. In: Proceedings of the 21st International Conference of Principles and Practice of Constraint Programming (CP 2015). Springer. 2015, pp. 592–608.

[4] Mathieu Jadin et al. “CG4SR: Near optimal traffic engineering for segment routing with column generation”. In: INFOCOM Conference on Computer Communications. IEEE. 2019, pp. 1333–1341.

Session 12B: Invited session: Discrete optimization for hydro unit commitment problem *Room: Henri Poincaré (Amphi I)*

Branch-and-Cut en 2 phases pour gérer les symétries polyédrales, application au 1-HUC

Pascale Bendotti, Alexandre Heintzmann, Cécile Rottner

In this paper, we consider polyhedral symmetries, defined as permu- tations of variables that preserve feasibility independently of the objec- tive function. Unlike classical symmetry-handling techniques, we develop a Branch&Cut framework with a separation in two-phases that leverages these symmetries to structure families of valid inequalities and enhance their separation. The first phase generates representative valid inequalities based solely on the constraint set, while the second phase separates their symme- try orbits during the Branch-and-Cut process. For column-wise symmetric polyhedral symmetry groups, we show that the orbit-separation problem re- duces to a maximum matching problem which can be solved in polynomial time. We apply the proposed framework to the Symmetric-weight Chain Precedence Knapsack problem, a knapsack variant arising in Hydro Unit Commitment problems at EDF. Computational results show that exploiting polyhedral symmetries significantly strengthens the separation process and improves overall solution performance.

Session 12K: Multiobjective Optimization and Evolutionary Computation Part I *Room: Loggia (A1 - first floor)*

A Fresh Look at Lamarckian Evolution and the Baldwin Effect

Inès Benito, Johannes F. Lutzeyer, Benjamin Doerr

LIX, CNRS, École Polytechnique, IP Paris, France, {firstname.lastname}@polytechnique.edu

Keywords: Memetic algorithms, Baldwin effect, runtime analysis

Baldwinian and Lamarckian evolution have existed for a long time in evolutionary algorithms (EAs) without ever dominating the academic literature or practical applications. In this work, we use modern empirical and theoretical methods to revisit Lamarckian and Baldwinian evolution and rigorously compare them with the generic Darwinian evolution. On the empirical side, we run a comprehensive suite of experiments on graphs from six different datasets from the recent GraphBench benchmark on Maximum Independent Set and Maximum Cut problems. Our results show that Baldwinian and Lamarckian evolution consistently outperform Darwinian evolution, confirming the great potential of local search augmented evolutionary algorithms. Notably, in the great majority of cases, all EAs outperform recent deep learning baselines and approach the performance of highly specialised heuristic and exact solvers. We furthermore report a high-performing set of generalist parameters for all studied evolution types that we hope will be of use to practitioners in future. On the theoretical side, we extend the existing DeceptiveLeadingBlocks benchmark to arbitrary block length k. For all constant k, we then prove asymptotically tight runtime bounds for the (1 + 1) EA in the three evolution types on this benchmark. For Baldwinian evolution, these are independent of k, whereas for the other two evolution types, the runtimes steeply increase with growing value of k.

References

[1] I. Benito, J. F. Lutzeyer, B. Doerr, A fresh look at Lamarckian evolution and the Baldwin effect, Parallel Problem Solving from Nature, PPSN 2026, Lecture Notes in Computer Science 16988, 459–474, Springer, 2026, doi:10.1007/978-3-032-36229-2_28.

Session 15J: Operations Research for Food and Healthcare *Room: Emmy Noether (A1.140)*

Lagrangian heuristic for a food lot-sizing problem

Lucas Berterottière[1], Adeline Fermanian[2], Safia Kedad-Sidhoum[3]

[1]Sorbonne Université, France, lucas.berterottiere@sorbonne-universite.fr
[2]Califrais, France, adeline.fermanian@califrais.fr
[3]CNAM, France, safia.kedad_sidhoum@cnam.fr

Keywords: Lot-sizing, Lagrangian relaxation, Dynamic programming, Food supply chain

Estimates suggest that about 14% of the food produced are wasted between production and distribution [2]. This waste leads to an economic cost for the companies involved and an ecological cost for the society. This abstract addresses an inventory problem with operational constraints from Califrais, the digital operator of Rungis, with a consideration on food waste. The problem studied is a multi-item uncapacitated lot-sizing problem where several constraints are considered: the minimum order quantity, the perishability of products, and the inventory bound. Each item has a unit volume, a lifespan and a minimum order quantity, and unsold units reaching their lifespan must be disposed of at a cost penalizing food waste. The items are coupled by the storage capacity of the warehouse. This problem is a harder version of the multi-item uncapacitated lot-sizing problem with inventory bound, proven strongly NP-hard in [1], so it is NP-hard as well. As the inventory bound constraint is the only coupling constraint, a Lagrangian relaxation heuristic is made to decompose the problem into simple single-item lot-sizing problems. We prove that the resulting single-item lot-sizing problems can be solved polynomially using dynamic programming. The subplan decomposition used to build the recurrence will be presented, as well as the dominance properties derived from the constraints. As the Lagrangian relaxation modifies production, holding and disposal costs, a cost transformation is made to obtain non-speculative disposal costs. A repair heuristic restores the storage capacity at each subgradient iteration, providing feasible solutions whose quality is assessed on instances derived from [3], as well as on real-case instances from Califrais. A performance comparison is given against mixed-integer linear programming formulations and sensitivity analysis of multiples parameters will be shown.

References

[1] A. Akbalik, B. Penz, C. Rapine, Multi-item uncapacitated lot sizing problem with inventory bounds, Optimization Letters, 9(1):143–154, 2015.

[2] FAO, The State of Food and Agriculture 2019. Moving forward on food loss and waste reduction, Rome, 2019.

[3] W. W. Trigeiro, L. J. Thomas, J. O. McClain, Capacitated lot sizing with setup times, Management Science, 35(3):353–366, 1989.

Session 12F: Routing & Mobility *Room: Ada Lovelace (A1.128)*

Pricing for Deliveries Pooling

Nicolas Bessin, Guillaume Dalle, Massil Hihat, Frédéric Meunier

1CERMICS, CNRS, ENPC, Institut Polytechnique de Paris, Marne-la-Vallée, France,firstname.lastname@enpc.fr; 2LVMT, Univ Gustave Eiffel, ENPC, Institut Polytechnique de Paris, Marne-la-Vallée, France, firstname.lastname@enpc.fr; 3Califrais, Paris, France, firstname.lastname@califrais.fr

Keywords: Pooling in Logistics, Collaborative Vehicle Routing, Cost Allocation

In the fresh-food supply chain—specifically at the Rungis food market, multiple suppliers each sell a different type of goods (meat; fish; and fruits and vegetables). Each must independently perform deliveries to their own set of clients, and there is heavy overlap between these sets of clients, most of them being located in close proximity to each other. This independent routing leads to two types of inefficiencies: partially filled trucks and multiple visits to the same client. To alleviate this, Califrais aims to introduce a new actor to the system, called the "pooling operator," that can take over some of the deliveries of the different suppliers and pool them, leading to better efficiency than what the suppliers can achieve on their own. We propose a new form of (vertical) collaborative routing model [1], where the pooling operator provides a "pooling service" that the suppliers can use to reduce their own routing costs. In our model, the pooling operator prices the service in advance and chooses daily how to operate the system. To make sure the different suppliers want to participate, the pooling operator must ensure that operating costs are reduced for suppliers who participate in the pooling scheme, while trying to maximize its own profit. Our main contribution lies in using a pooling model where the allocation of profits is done in advance rather than through the use of a posteriori reallocation of the profits, which is the more common approach in the literature [2]. We formulate the model, provide bounds on the usefulness of pooling in simple cases, and state properties that will help us solve the model. Finally, we provide initial numerical results that demonstrate the viability of the proposed pooling system, based on a real-world data set from suppliers operating at the Rungis food market.

References

[1] Margaretha Gansterer and Richard F. Hartl. "Collaborative Vehicle Routing: A Survey". In: European Journal of Operational Research 268.1 (July 2018), pp. 1–12. issn: 0377-2217. doi: 10.1016/j.ejor.2017.10.023. (Visited on 04/30/2026).

[2] Mario Guajardo and Mikael Rönnqvist. "A Review on Cost Allocation Methods in Collaborative Transportation". In: International Transactions in Operational Research 23.3 (May 2016), pp. 371–392. issn: 0969-6016, 1475-3995. doi: 10.1111/itor.12205. (Visited on 04/22/2026).

Session 12I: Optimization for Learning and Large-Scale Decision Making *Room: Sophie Germain (A1.139)*

Consensus-Based Optimization Beyond Finite-Time Analysis

Pascal Bianchi, Radu-Alexandru Dragomir, Victor Priser

LTCI, Télécom Paris, France

We analyze a zeroth-order particle algorithm for the global optimization of a non-convex function, focusing on a variant of Consensus-Based Optimization (CBO) with small but fixed noise intensity. Unlike most previous studies restricted to finite horizons, we investigate its long-time behavior with fixed parameters. In the mean-field limit, a quantitative Laplace principle shows exponential convergence to a neighborhood of the minimizer x_*. For finitely many particles, a analysis based on propagation of chaos yields explicit error bounds: individual particles achieve long-time consistency near x_*. We also analyze the performance of the "global best" particle. The proof technique combines a quantitative Laplace principle with block-wise control of Wasserstein distances, avoiding the exponential blow-up typical of Grönwall-based estimates.

Session 15H: Invited Session:Neural Network approximations for high-dimensional PDEs Problems *Room: Alan Turing (A1.134)*

Solving Hamilton-Jacobi equations by residual minimization of monotone finite-difference discretizations

Olivier Bokanowski, Carlos Esteve-Yagüe, Richard Tsai

1Université Paris Cité, LJLL, France, olivier.bokanowski@u-paris.fr 2Universidad de Alicante, Spain, carlos.esteve@ua.es 3Oden Institute, University of Texas at Austin, USA, ytsai@math.utexas.edu

Keywords: Hamilton-Jacobi equations, residual minimization, monotone schemes, neural networks, viscosity solutions, Polyak–Łojasiewicz inequality

We introduce a method for solving Hamilton–Jacobi equations, both inviscid and viscous, by minimizing the squared residuals of monotone finite-difference discretizations on grids of varying resolution. In contrast to the uniform-grid setting of [1], our analysis and numerical framework accommodate possibly non-uniform grids. The method is designed to leverage neural networks and modern GPUs to solve these equations in higher dimensions; consequently, the setting for our analysis is the minimization of the residual functionals via gradient-based optimization. We establish a well-posedness theory for this approach: any critical point of the finite-difference loss solves the monotone scheme together with the prescribed Dirichlet boundary conditions, and the error of an approximation is controlled by its residual. We then derive the rate of convergence of the gradient flow that minimizes the residual in several settings, noting that the rate may depend on the grid resolution and the domain dimension. Building on this foundation, we propose a multi-level training algorithm that exploits the faster convergence available on the coarser grids of the discretization. Combined with the convergence theorem of Barles and Souganidis for monotone and consistent schemes, our results guarantee convergence to the unique viscosity solution as the grid is refined. We illustrate the approach on eikonal equations, level-set problems, and a Hamilton–Jacobi–Isaacs equation arising from a stochastic differential game, in dimensions up to eight [2].

References

[1] C. Esteve-Yagüe, R. Tsai, and A. Massucco. Finite-difference least square methods for solving Hamilton-Jacobi equations using neural networks. Journal of Computational Physics, 524:113721, 2025. https://arxiv.org/abs/2406.10758

[2] C. Esteve-Yagüe, O. Bokanowski, and R. Tsai. Solving Hamilton-Jacobi equations by residual minimization of monotone finite-difference discretizations. Preprint, 2026. https://arxiv.org/abs/ 2601.21764

Session 15H: Invited Session:Neural Network approximations for high-dimensional PDEs Problems *Room: Alan Turing (A1.134)*

Neural Feedback Learning for Stochastic Control with Degenerate Diffusions: Quantitative Error Estimates and Numerical Experiments

Olivier Bokanowski, Jean-François Chassagneux, Marco Scaratti, Xavier Warin

1Université Paris Cité, Laboratoire Jacques-Louis Lions, France, olivier.bokanowski@u-paris.fr 2ENSAE, CREST and Institut Polytechnique de Paris, France, jean-francois.chassagneux@ensae.fr 3University of Verona, Italy, marco.scaratti@univr.it 4EDF Lab Paris-Saclay and FiME, France, xavier.warin@edf.fr; Parallel Session: Neural Network approximations and related techniques for high-dimensional PDEs and Control Problems.

Keywords: Stochastic optimal control, degenerate diffusions, neural feedback control, error estimates.

We study finite-horizon stochastic optimal control problems and approximate the resulting time-discrete formulation by a direct policy-learning problem over neural-network feedback maps. We prove a quantitative convergence estimate, in an averaged sense, for the error between the time-discrete value and the value induced by an approximately optimized neural policy. The bound separates the approximation of near-optimal feedback policies, the localization of stochastic trajectories on compact sets, and the optimization tolerance in training. The analysis does not require transition-density assumptions and covers possibly degenerate diffusions and deterministic controlled dynamics in a unified framework. Numerical experiments are provided for a degenerate stochastic radial target problem, a Hamilton Jacobi Bellman benchmark, and a gas storage problem, illustrating the approach and separating the main error sources: time discretization, restriction to piecewise-constant policies, neural-network approximation, and Monte Carlo evaluation.

References

[1] O. Bokanowski, J.-F. Chassagneux, M. Scaratti, X. Warin, Neural feedback approximation for stochastic control with degenerate diffusions: error estimates and numerical analysis, arXiv:2607.10254.

Session 15A: Optimal Control, Measures and PDE constrained optimization *Room: Marie Curie (Auditorium)*

The Adjoint State Method for Parametric Definable Optimization without Smoothness or Uniqueness

Jérôme Bolte[1], Edouard Pauwels[1], Cheik Traoré[1]

[1]Toulouse School of Economics, Toulouse, France, cheik.traore@tse-fr.eu

Keywords: Adjoint state method, conservative fields, nonsmooth optimization, conic programming, o-minimal geometry, parametric optimization, variational analysis.

We establish that nonconvex definable parametric optimization problems with possibly nonsmooth objectives, inequality constraints, conic constraint systems, and non-unique primal and dual solutions admit an adjoint state formula under a mere qualification condition. The adjoint construction yields a selection of a conservative field for the value function, providing a computable first-order object without requiring differentiation of the solution mapping. Through examples, we show that even in smooth problems, the formal adjoint construction fails without conservativity or definability, illustrating the relevance of these concepts to grasp theoretical aspects of the method. This work provides a tool that can be directly combined with existing primal-dual solvers for a wide range of parametric optimization problems.

References

[1] J. Bolte, E. Pauwels, C. Traoré, The adjoint state method for parametric definable optimization without smoothness or uniqueness, arXiv:2603.26503.

Session 8H: Learning dynamics and stochasticity in optimization *Room: Alan Turing (A1.134)*

Convergence of Optimizers Implies Eigenvalues Filtering at Equilibrium

Jérôme Bolte, Quoc-Tung Le, Edouard Pauwels

1Toulouse School of Economics, Universite Toulouse Capitole, France, jerome.bolte@tse-fr.eu, edouard.pauwels@tse-fr.eu
2Univ. Grenoble Alpes, CNRS, LJK, Grenoble, France, quoc-tung.le@univ-grenoble-alpes.fr

Keywords: Implicit bias, dynamical systems, Hadamard–Perron theorem, sharpness-aware minimization

Ample empirical evidence in deep neural network training suggests that a variety of optimizers tend to find nearly global optima. In this work, we adopt the reversed perspective that convergence to an arbitrary point is assumed rather than proven, focusing on the consequences of this assumption. From this viewpoint, in line with recent advances on the edge-of-stability phenomenon, we argue that different optimizers effectively act as eigenvalue filters determined by their hyperparameters. Specifically, standard gradient descent inherently avoids the sharpest minima, whereas Sharpness-Aware Minimization (SAM) algorithms go even further by actively favoring wider basins. Inspired by these insights, we establish a generalized Hadamard–Perron stable manifold theorem applying to general semialgebraic C^2 functions, without requiring additional non-degeneracy conditions or global Lipschitz bound assumptions. We leverage this theoretical framework to propose two novel algorithms—Two-step USAM and Hessian USAM—that exhibit enhanced eigenvalue filtering and promote wider minima. Numerical experiments on feed-forward neural networks and ResNet architectures support our theoretical conclusions.

References

[1] Morris Hirsch, Charles Pugh, and Michael Shub. Invariant Manifolds, volume 76. Springer Berlin, Heidelberg, 1977

[2] Jason D. Lee, Max Simchowitz, Michael I. Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir (eds.), 29th Annual Conference on Learning Theory, 2016

[3] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur, Sharpness-aware minimization for efficiently improving generalization, In International Conference on Learning Representations (ICLR), 2021.

Session 5J: Invited session : Quantum Applications in Operations Research part 1 *Room: Emmy Noether (A1.140)*

An alternative exploration of the solutions space for the Travelling Salesman Problem leveraging a warm-start and a preserving mixer to extend a recent proposal

Martin Bombardelli, Philippe Lacomme

[1]Université Clermont Auvergne, LIMOS; [2]Université Clermont Auvergne, LPCA; [3]CGI, Clermont-Ferrand, France; [4]EURO WG on Quantum Computing

The travelling salesman problem (TSP) is one of the most studied problems in operations research. Many works focus on this problem both in classical and quantum optimization. For a N city TSP, we propose an improvement over the work of (Baldazzi et al., 2026) which came up with a 2 log2(N). Our proposal consists in a new mixer allowing us to put into superposition only valid TSP tours at the cost of an ancilla qubit. As in the original work, the resolution of the problem is through a variational quantum circuit (VQA). Examples with five and eight cities are solved and two different methods to extract the solutions are given.

Theoretical circuit

The proposed circuit keeps the structure of the original work but replaces every step. A warm start obtained through classical heuristics is used to prepare an initial state. The initial state preparation involves the state preparation of (Shukla et al.,2023) and the solution found by classical means.
Then, an operator to exchange two different city amplitudes on a register (BSWAP) is introduced. Using this operator, one can build another operator (Wij) which will exchange two cities i and j in a tour in the 2 log2 N representation.

Solution extraction and results

We propose a procedure to extract the solutions from optimized angles. This involves rounding the angles obtained and applying or not the exchanges made in the quantum circuit onto the warm start

$$\forall k \in \left\{0, \dots, \frac{(N-1)(N-2)}{2} - 1\right\}, \qquad \theta_k^{\text{opt}} = \begin{cases} 0, & \theta \leq \pi/4, \\ \pi/2, & \theta > \pi/4. \end{cases}$$

Experiments were run on a simulator using Qiskit. One example of result is given in Table 1 for a 5- cities instance where the arcs included in the optimal solution are found with probability 0.94.
Solution Cost Probability to appear
[0, 3, 1, 2, 4, 0] 18 0.94
[0, 1, 4, 2, 3, 0] 219 0.03
[0, 3, 1, 4, 2, 0] 144 0.02
For the 8-city example, the solution extracted from the angles was not optimal but, using the X matrix and the extraction from (Baldazzi et al., 2026), the optimal solution was still recovered.

References

Baldazzi, A., Azzini, S., & Pavesi, L. (2026). Resource-efficient variational quantum solver for the traveling salesman problem and its silicon photonics implementation. Quantum Science and Technology, 11(3), 035057. doi:10.1088/2058-9565/ae8df2.
Shukla, A., & Vedula, P. (2024). An efficient quantum algorithm for preparation of uniform quantum superposition states. Quantum Information Processing, 23(2), 38. doi:10.1007/s11128-024-04258-4.

Session 12D: Optimal Transport & Distributional Robustness *Room: Lise Meitner (A1.116)*

Quadratically Regularised Wasserstein Medians

Benoît Bonnet-Weill, Matteo Tacchi-Bénard, Nicola Zaupa

1Université Paris-Saclay, CNRS, L2S CentraleSupélec benoit.bonnet-weill@centralesupelec.fr 2Université Grenoble Alpes, Grenoble INP, CNRS, GIPSA-lab. matteo.tacchi@grenoble-inp.fr, nicola.zaupa@gipsa-lab.fr

Keywords: Wasserstein Aggregators, Breakdown Point, Quadratic Regularisation

In this work, we consider quadratically regularised medians in Wasserstein spaces and investigate their stability properties with respect to adversarial inputs. This model is a measuretheoretic version of that recently proposed in [?] in the context of Bayesian decision making over Euclidean spaces. In the latter work, it shown among other things that adding inside a median problem a Tikhonoff term penalising the distance from a prior belief entails strong resilience properties on the resulting aggregator, wherein the total displacement induced by an adversary input is bounded by the ratio between its weight and the regularisation parameter, regardless of its spatial location. The corresponding notion of Lipschitz resilience is currently used for robust content recommendations on the platform Tournesol.app, through a bayesian generalization of Bradley-Terry models, see e.g. [?]. After measuring the sensitivity to adversaries of classical Wasserstein aggregators in terms of their breakdown points heuristically, the minimal amount of inputs that one need to manipulate to make the aggregator diverge , we introduce two different Wasserstein median models which are amenable to quadratic regularisation. The first one involves linearised W2-metric at the prior, and leverages convexity properties along generalised geodesics, while the second is based on the vanilla W1-median investigated in [?], and hinges upon fine flat uniform convexity estimates for the squared W2-metric explored in [?]. The improved robustness of these new aggregators is illustrated via numerical simulations.

References

[1] Y. Allouah, R. Guerraoui, L.-N. Hoang, and O. Villemaud. Robust Sparse Voting. In International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, 2024.

[2] A. Delalande and Q. Merigot. Quantitative Stability of Optimal Transport Maps under Variations of the Target Measure. Duke Mathematical Journal, 172(17):3321-3357, 2023.

[3] G. Carlier, E. Chenchene, and K. Eichinger. Wasserstein Medians: Robustness, PDE Characterization, and Numerics. SIAM Journal on Mathematical Analysis, 56(5):6483-6520, 2024.

[4] S. Fageot, S. Farhadkhani, L.-N. Hoang, and O. Villemaud. Generalized Bradley-Terry Models for Score Estimation from Paired Comparisons. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 20379-20386, 2024.

Session 5K: Decision Focused Learning *Room: Loggia (A1 - first floor)*

Integrated Model for Optimization and Machine Learning: the Case of Predictive Maintenance

Marie Bouilloud[1], Michel Nakhla[1]

[1]École des Mines Paris-PSL, France, marie.bouilloud@minesparis.psl.eu, michel.nakhla@minesparis.psl.eu

Keywords: Decision-making under uncertainty, Optimization, Mixed Integer Linear Programming, Predictive maintenance

Decision-making under uncertainty involves making a decision at time where all of the elements are not known. If, instead, contextual data is available, it is possible, with data-driven methods, to predict the missing elements. There are thus two challenges: prediction, using machine learning, and decision, using constrained optimization. The predict-then-optimize sequential paradigm proposes to predict the missing elements, with the machine learning model aiming at minimizing the error between the predicted and the true parameters. However this method has been proven to be structurally sub-optimal [1]. Decision-focused learning proposes that the machine learning model learns to minimize a function of regret over the final decision, under strong differentiability conditions [2]. Predictive maintenance is one of these cases, with a failure prediction problem and a maintenance planning optimization problem. The maintenance planning optimization problem is modeled as a mixed-integer linear program aiming at minimizing the total maintenance cost under resources and potentiality constraints. Decision-focused learning is not applicable in this case because the uncertainty concerns the set of decision variables. The set of tasks to plan is indeed not fixed. Furthermore, there is an intermediary step between prediction and decision consisting in transforming the failure predictions into predictive tasks. This paper proposes a model integrating prediction and optimization to serve a decision-aid tool for farm operators. The integrated model is original and integrates a function of regret over the decision as the score for the hyperparameters selection of the machine learning model. Experiments are carried out based on data from industrial partners, namely wind and solar production assets operators, and show that, on the tested instances, a predictive maintenance model allows to avoid curative maintenance. Furthermore, the proposed integrated model reduces over-maintenance over a sequential model.

References

[1] A. N. Elmachtoub and P. Grigas, Smart "Predict, then Optimize", Management Science, 68(1):9–26, 2022.

[2] J. Mandi, J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns and F. Fioretto, Decision-focused learning: Foundations, state of the art, benchmark and future opportunities, Journal of Artificial Intelligence Research, 80, 1623-1701, 2024.

Session 12G: Mid and Long Term Energy Management *Room: Enrico Fermi (A1.133)*

Distributed Optimization for Hydro-thermal Scheduling using ADMM Framework

Matias Bresciano, Hamid Ben Ahmed, Roman Le Goff Latimier, Jean-Marc Janin, Juliette Gerbaux

1RTE Réseau de Transport d'Electricité, France, matias.bresciano@rte-france.com, jean-marc.janin@rte-france.com, juliette.gerbaux@rte-france.com 2Université de Rennes, ENS Rennes, IETR, CNRS, roman.legoff-latimier@ens-rennes.fr, benahmed@ens-rennes.fr

Keywords: Hydro-thermal Scheduling, Multi-stage Stochastic Optimization Problem, Distributed Optimization, Alternating Direction Method of Multipliers, Dual Approximate Dynamic Programming

Optimizing interconnected hydro-thermal systems raises two structural difficulties, in particular with multi-reservoir prospective studies led by RTE [1]. First, dealing with the uncertainties requires a multi-stage stochastic optimization formulation. Second, spatial coupling through interconnections prevents each node from being treated independently. Classical algorithms such as Stochastic Dynamic Programming (SDP) rapidly face the curse of dimensionality when the numbers of reservoirs to optimize rise [2]. Price decomposition schemes such as Dual Approximate Dynamic Programming (DADP) decouple the problem spatially, resulting in a non-smooth dual problem. Existing works solve it either with subgradient ascent, whose convergence requires diminishing step-size rules and is slow in practice, or with L-BFGS, whose guarantees assume differentiability and therefore do not apply here [3]. Moreover, its convergence analysis relies on convexity, which in practice is obtained by assuming linear reservoir dynamics, a modeling approximation we aim to relax. In order to tackle these issues, a new framework is presented, based on the Alternating Direction Method of Multipliers (ADMM) [4]. This framework enables convergence in linear and non-linear cases. Unlike the DADP framework, which implements deterministic price decomposition processes, we aim to implement stochastic dual multipliers, which are assumed to follow a Markovian process. The dynamic of the dual multiplier is learned through the process [5]. Results obtained are compared to exact reference obtained with SDP and approximated result obtained with DADP. This work paves the way to optimizing multi-reservoir systems in a disaggregated fashion, which represents a significant achievement for transmission system operators such as RTE.

References

[1] RTE, Futurs Energétiques 2050 : Rapport Complet (2021).

[2] Bellman, R. (1966). Dynamic programming. science, 153(3731), 34-37.

[3] Parra, C. M. (2025). Valeurs d'usage dans un système multinœuds pour des études prospectives en énergie (Doctoral dissertation, École des Ponts ParisTech).

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends in Machine Learning, 3(1), 1–122.

[5] R. L. G. Latimier, G. Chérot, H. B. Ahmed (2022). Online learning for distributed optimal control of an electric vehicle fleet. Electric Power Systems Research, 212, 108330.

Session 8J: Optimal Tariffs *Room: Emmy Noether (A1.140)*

Joint Design of Time-of-Use Schedules and Price Levels: a Bilevel Model with Client Participation and Load Flexibility

Abdellah Bulaich Mehamdi, Wim Van Ackooij, Luce Brotcorne, Stephane Gaubert, Quentin Jacquet

[1]CMAP, Ecole polytechnique, Institut Polytechnique de Paris, Palaiseau, France; [2]EDF Lab Paris-Saclay, OSIRIS department, Palaiseau, France; [3]Inria, France

Keywords: Bilevel optimization; Time-of-use pricing; Demand response; Electricity retail; Nonsmooth optimization; Stochastic gradient ascent; Mixed-integer nonlinear programming

We study the joint design of a time-of-use (ToU) tariff (the partition of the day into pricing blocks and the price level attached to each block) by a profit-maximizing electricity retailer facing flexible clients who may refuse the offer in favour of an outside option. The problem is modeled as a bilevel program in which each client, at the lower level, reshapes its consumption under pointwise bounds and a fixed daily energy total, and simultaneously chooses a participa- tion level penalized by a quadratic regularizer. We prove the client's best response is computable in loglinear time. We derive two equivalent single-level reformulations: a closed-form continuous one tackled via stochastic gradient ascent (using a continuous reparametrization of block start times and durations) and a mixed-integer nonlinear program solved via dedicated solvers. The stochastic ascent scheme provides Clarke-critical accumulation points almost surely, in a scalable way. On a case study from the French retail electricity market with 5000 simulated households, the gradient method matches a commercial mixed-integer solver on the pointwise problem to within 0.3% of the optimal profit while running two to eight times faster on 100 representative clients. Restricting unconstrained pointwise prices to admissible ToU schedules yields an esti- mated profit loss of 1.3%-13.8% for two-period tariffs and 0.6%-6.6% for three-period tariffs on the instances considered.

Session 15E: Algorithms 2.0 : advances in computational methods *Room: Antoine Lavoisier (A1.122)*

A New Bi-Objective View of Bundle Methods

Enrico Calandrini[1], Antonio Frangioni[2]

[1]Department of Energy, Systems, Territory and Construction Engineering, University of Pisa, Italy, e.calandrini@studenti.unipi.it
[2]Department of Computer Science, University of Pisa, Italy, antonio.frangioni@unipi.it

Keywords: nonsmooth convex optimization; bundle methods; bi-objective optimization; merit function

We propose a new perspective on bundle methods that, rather than relying on the classical combination of a cutting-plane lower model and a stabilization mechanism [4, 5, 3], models the objective simultaneously from below and above. In this framework, a lower model fB and an upper model fP jointly define the bi-objective problem minx (fB(x), fP (x)) , which provides the master-problem formulation of the method. We show that, through appropriate choices of the two models, classical stabilization techniques - including proximal, level, and trust-region methods [5, 6, 8] - can be interpreted as different rules for selecting points on the same Pareto frontier generated by (fB, fP). In the polyhedral setting, this viewpoint connects naturally with the framework of Bertsekas and Yu [1], where outer and inner linearizations play roles analogous to our lower and upper models, respectively. A merit function is then introduced to measure both model agreement and approximate optimality. Rather than prescribing the serious/null decision and the aggregation of cutting planes externally, we determine them by directly minimizing this merit function. In a two-point aggregation scheme, this leads to an explicit update: the center decision emerges as a binary outcome of merit minimization, while the optimal cut-aggregation parameter is obtained in closed form. Numerical experiments indicate that the resulting Upper–Lower Model Bundle Method is competitive with the classical proximal bundle method and that the associated model gap provides an effective proxy for true accuracy. We also prove convergence under the actual merit-minimizing update rule and establish an $O(1/\varepsilon 2)$ complexity bound for the associated optimality certificate. We then show that, by combining bundle refinement with a relativeaccuracy serious/null step mechanism and a proximal-distance argument, the framework recovers the optimal $O((LD/\varepsilon)2)$ nonsmooth first-order complexity bound for true objective accuracy [7, 2].

References

[1] D. P. Bertsekas and Y. Huizhen. A unifying polyhedral approximation framework for convex optimization. SIAM Journal on Optimization, 21(1):333–360, 2011.

[2] M. Díaz and B. Grimmer. Optimal convergence rates for the proximal bundle method. SIAM Journal on Optimization, 33(2):424–454, 2023.

[3] A. Frangioni. Generalized bundle methods. SIAM Journal on Optimization, 13(1):117–156, 2002.

[4] J. E. Kelley, Jr. The cutting-plane method for solving convex programs. Journal of the Society for Industrial and Applied Mathematics, 8(4):703–712, 1960.

[5] C. Lemaréchal. Bundle methods in nonsmooth optimization. In Claude Lemaréchal and Robert Mifflin, editors, Nonsmooth Optimization, volume 3 of IIASA Proceedings Series, pages 79–102. Pergamon Press, Oxford, 1978.

[6] C. Lemaréchal, A. Nemirovski, and Y. Nesterov. New variants of bundle methods. Mathematical Programming, 69:111–147, 1995.

[7] A. S. Nemirovski and D. B. Yudin. Problem Complexity and Method Efficiency in Optimization. John Wiley & Sons, Chichester, 1983.

[8] H. Schramm and J. Zowe. A combination of the bundle approach and the trust region concept. In Advances in Mathematical Optimization. Walter de Gruyter, Berlin, 1988.

Session 15E: Algorithms 2.0 : advances in computational methods *Room: Antoine Lavoisier (A1.122)*

BundleSolver 2.0: A Modular Implementation of Generalized Bundle Methods

Enrico Calandrini[1], Donato Meoli[2], Antonio Frangioni[2]

[1]Department of Energy, Systems, Territory and Construction Engineering,
University of Pisa, Italy, e.calandrini@studenti.unipi.it
[2]Department of Computer Science, University of Pisa, Italy,
d.meoli@studenti.unipi.it, antonio.frangioni@unipi.it

Keywords: bundle methods; nonsmooth convex optimization; mathematical decomposition

Bundle methods are a versatile framework for nonsmooth convex optimization, particularly for structured problems arising from decomposition and Lagrangian relaxation [3]. However, their effectiveness relies on how implementations exploit problem structure, manage accumulated information, and interface with master-problem solvers [4, 1]. We present BundleSolver 2.0 [2], a general-purpose implementation of generalized bundle methods for nonsmooth convex optimization within the SMS++ framework [6]. Although the solver does not require a decomposable objective, SMS++ naturally supports structured problems and motivates a specialized treatment of objectives expressed as sums of components. In this setting, components that are sufficiently "easy" can be represented exactly in the master problem [5], while each remaining component is approximated independently through a cutting-plane model constructed from first-order oracle information. A central design principle of BundleSolver 2.0 is the separation between the bundle algorithm, the representation of the master problem, and its numerical solution [2]. A dedicated master-problem component encapsulates the cuttingplane models and stabilization terms, while numerical optimization is delegated to solvers through the modeling infrastructure provided by SMS++ [6]. This architecture supports both primal and dual master formulations within a unified framework and accommodates proximal, level, and doubly stabilized bundle variants [2]. We discuss key implementation choices, including incremental master-problem updates [7, 2], linearization management and aggregation, sparse componentvariable dependencies, and efficient handling of problem modifications [2]. The goal is to combine flexibility and maintainability with computational efficiency, while supporting experimentation with alternative master formulations and stabilization strategies.

Finally, we present preliminary numerical results on selected applications, illustrating the computational effectiveness of the proposed implementation and the impact of different design choices on the behavior and performance of the solver.

References

[1] L. Bacaud, C. Lemaréchal, A. Renaud, and C. Sagastizábal. Bundle methods in stochastic optimal power management: A disaggregated approach using preconditioners. Computational Optimization and Applications, 20(3):227–244, 2001.

[2] E. Calandrini, D. Meoli, and A. Frangioni. BundleSolver 2.0: A generalized bundle solver for SMS++. GitLab repository, BundleSolver2.0 branch, 2026. SMS++ project.

[3] A. Frangioni. Generalized bundle methods. SIAM Journal on Optimization, 13(1):117–156, 2002.

[4] A. Frangioni and G. Gallo. A bundle type dual-ascent approach to linear multicommodity min-cost flow problems. INFORMS Journal on Computing, 11(4):370–393, 1999.

[5] A. Frangioni and E. Gorgone. Generalized bundle methods for sum-functions with "easy" components: Applications to multicommodity network design. Mathematical Programming, 145(1–2):133–161, 2014.

[6] SMS++ Project. SMS++: A C++ library for modeling and solving block-structured mathematical models. GitLab repository, 2026.

[7] W. van Ackooij and A. Frangioni. Incremental bundle methods using upper models. SIAM Journal on Optimization, 28(1):379–410, 2018.

Session 5C: IS: Modelling supply chains in Mathematical Programming: flow and economics *Room: André Marie Ampère (Amphi II)*

A mathematical programming formulation for generalized supply chains

Justin Capgras, Haitam Chakori, Claudia D'Ambrosio, Michele Fabi, Leo Liberti

1 Digital Excellence Center, Assystem, France, jcapgras@assystem.com 2LIX CNRS, École Polytechnique, Institut Polytechnique de Paris, F-91128 Palaiseau, France, {chakorisemmane,dambrosio,liberti}@lix.polytechnique.fr 3 Télécom Paris, Palaiseau, France, michele.fabi@telecom-paris.fr

Keywords: Supply chain, Mathematical programming, Optimization, Generalized flow

This talk is part of the invited session Modelling supply chains in Mathematical Programming: flow and economics (organizer: Leo Liberti)

Supply-chain optimization has generated a very large literature, but most mathematical models focus on a specific combination of decisions rather than on a unified representation of the complete system [1]. These approaches considerably extend classical supply-chain models. Some authors explicitly introduce the term generalized supply chain. However, "generalize" does not necessarily have the same meaning across the literature. In many cases, generalization concerns the topology of the network, the type of facilities considered, or the integration of several predefined decisions. The underlying production structure generally remains applicationdependent [2]. The approach investigated here adopts a different perspective. We seek to develop a formulation in which transportation, storage and arbitrary transformations of commodities are primitive mathematical operations. A transformation is not restricted to a predefined manufacturing process or to a mixture of materials: a commodity h entering a plant may generate another commodity k, possibly with a yield different from one and with an associated processing time. This unified framework provides a basis for modeling a broad range of integrated supplychain systems and may facilitate future extensions incorporating additional operational decisions and application-specific constraints representing production, assembly, conversion, dismantling or disposal through the same mathematical object. We present a mathematical programming formulation for this object and some extensions which allow modeling well-known supply chain problems such as multicommodity flow, blending problem, scheduling, etc.

References

[1] Dušan Hrabec, Lars Magnus Hvattum, Arild Hoff, The value of integrated planning for production, inventory, and routing decisions: A systematic review and meta-analysis, International Journal of Production Economics, Volume 248, 2022, 108468, ISSN 0925-5273

[2] Leo Liberti, Optimization and sustainable development. Computational Management Science, 12(3):371–395, 2015.

Session 5C: IS: Modelling supply chains in Mathematical Programming: flow and economics *Room: André Marie Ampère (Amphi II)*

A mathematical programming formulation for generalized supply chains - Industrial and Economic mechanisms

Justin Capgras, Haitam Chakori, Claudia D'Ambrosio, Michele Fabi, Leo Liberti

1 Digital Excellence Center, Assystem, France, jcapgras@assystem.com 2LIX CNRS, École Polytechnique, Institut Polytechnique de Paris, F-91128 Palaiseau, France, {chakorisemmane,dambrosio,liberti}@lix.polytechnique.fr 3 Télécom Paris, Palaiseau, France, michele.fabi@telecom-paris.fr

Keywords: Supply chain, Mathematical programming, Optimization, Reverse Supply Chain, Economic mechanisms, Generalized flow, Operations Management This talk is part of the invited session Modelling supply chains in Mathematical Programming: flow and economics (organizer: Leo Liberti)

Mathematical optimization has been used to model several parts of the supply chain, from production and inventory to transportation, reverse flows and environmental or operational constraints. Some models already combine several of these dimensions [2, 4].

However, even relatively generic formulation often remain tied to a predefined network structure [2]. We instead investigate whether different supply-chain structures can be represented through the same set of mathematical primitives, without redefining the formulation for each industrial configuration.

Our formulation builds on the transformation-flow concept introduced in [3]. The physical supply chain is described through transportation, storage and commodity transformation. We derive a compact formulation and an equivalent state-expanded network representation, whose structure can be exploited through a decomposition of the problem.

In this presentation, we focus on how this formulation can be extended from the physical description of the supply chain to its economic and industrial operation. We therefore introduce a set of operational mechanisms which are combined within a economic objective based on the value generated by the supply chain and the costs associated with its operation [1].

We then study whether introducing these mechanisms changes the mathematical structure identified in the original formulation. Several mechanisms can be incorporated while preserving the main structure of the formulation, whereas others introduce additional coupling or integer decisions and require a different treatment.

References

[1] Sunil Chopra. Supply Chain Management: Strategy, Planning, and Operation. Pearson, 8 edition, 2026. ISBN 9780135350126.

[2] Magdalini A. Kalaitzidou, Pantelis Longinidis, Panagiotis Tsiakis, and Michael C. Georgiadis. Optimal design of multiechelon supply chain networks with generalized production and warehousing nodes. Industrial & Engineering Chemistry Research, 53(33):13125–13138, 2014.

[3] Leo Liberti. Optimization and sustainable development. Computational Management Science, 12(3):371–395, 2015. doi: 10.1007/s10287-014-0219-5.

[4] Alyssa Mosca, Navneet Vidyarthi, and Ahmet Satir. Integrated transportation–inventory models: A review. Operations Research Perspectives, 6:100101, 2019.

Session 12H: Market Design and Strategic Behavior *Room: Alan Turing (A1.134)*

Trading with market resistance and concave price impact

Nathan De Carvalho[1], Youssef Ouazzani Chahdi[2], Grégoire Szymanski[3]

[1]LPSM, Université Paris Cité, decarvalho@lpsm.paris
[2]MICS, CentraleSupélec, youssef.ouazzani-chahdi@centralesupelec.fr
[3]DMATH, Université du Luxembourg, gregoire.szymanski@uni.lu

February 5, 2026

Keywords: Optimal trading, market resistance, concave market impact, propagator model, power-law decay, square root-law, Fredholm equations.

We consider an optimal trading problem under a market impact model with endogenous market resistance generated by a sophisticated trader who (partially) detects metaorders and trades against them to exploit price overreactions induced by the order flow. The model features a concave transient impact driven by a power-law propagator with a resistance term responding to the trader's rate via a fixed-point equation involving a general resistance function. We derive a (non)linear stochastic Fredholm equation as the first-order optimality condition satisfied by optimal trading strategies. Existence and uniqueness of the optimal control are established when the resistance function is linear, and an existence result is obtained when it is strictly convex using coercivity and weak lower semicontinuity of the associated profit-and-loss functional. We also propose an iterative scheme to solve the nonlinear stochastic Fredholm equation and prove an exponential convergence rate. Numerical experiments confirm this behavior and illustrate optimal round-trip strategies under "buy" signals with various decay profiles and different market resistance specifications.

Session 8G: Optimizing under network structure *Room: Enrico Fermi (A1.133)*

Optimal Power flow in Transmission Networks: a Spectral Graph Theory approach

G. Cavalcante, L. Melo, A. Knippel, J.-G. Caputo, R. Andrade

1 Departamento de Estatistica e Matematica Aplicada, Universidade Federal do Ceara, Brazil, geovanecavalcante@alu.ufc.br; rafael.andrade@ufc.br 2 Laboratoire de Mathématiques, INSA Rouen Normandie, France jean-guy.caputo@insa-rouen.fr; arnaud.knippel@insa-rouen.fr 3 Departamento de Engenharia Eletrica, Universidade Federal do Ceara, Brazil lucassmelo@dee.ufc.br

Keywords: Optimal power flow, power transmission network, spectral graph theory

This work proposes to use the spectral properties of the graph Laplacian matrix to balance load and optimize power flows in transmission lines. We develop two models: a quadratic formulation that minimizes the squared norm of line flows, solved efficiently through KarushKuhn-Tucker (KKT) conditions, and a linear formulation that directly minimizes the maximum line power across the entire system, promoting more balanced line loading. Computational experiments on Pandapower-adapted benchmark networks, including a large-scale 2848-bus system, demonstrate reductions of up to 39% in maximum line power, along with significant computational efficiency gains. The linear formulation can be naturally extended to incorporate realistic operational constraints, including line capacity limits, generation bounds, and phase angle restrictions. These results highlight the potential of spectral methods as an effective tool for enhancing load balancing, operational security, and resilience in modern power transmission networks.

References

[1] J.-G. Caputo, A. Knippel, and N. Retiere, Spectral analysis of load flow equations for transmission networks, Engineering Research Express, vol. 1, no. 2, p. 025007, 2019.

Session 5D: Invited session: Black-box Optimization *Room: Lise Meitner (A1.116)*

Black-box adversarial attacks on SPD Networks

Bastien Cavarretta, Thibault de Surrel, Clément W. Royer, Florian Yger

Université Paris Dauphine-PSL; Université Gustave Eiffel; INSA Rouen

Keywords: Black-box optimization, SPD Networks, Adversarial Attacks

This is an invited talk in the session "black-box optimization".

SPDNets are neural networks whose inputs lie on the manifold of symmetric positive definite matrices. Those networks are particularly useful on Brain-Computer Interface (BCI) tasks. However, they can be fooled by adversarial attacks based on Riemannian projected gradient descent, though designing such attacks requires complete, or white-box, access to the network. In this talk, we propose adversarial attacks on SPDNets assuming that the network is only available in a black-box fashion. Our attacks are based on direct-search and gradient-free algorithms, and, we leverage recent advances in Riemannian black-box optimization to tailor these methods to the Riemannian setting. Our experiments on BCI data show that our methods successfully fool the networks.

Session 5B: Invited session: Recent Advances in Mean Field Games and Mean Field Control *Room: Henri Poincaré (Amphi I)*

Learning controls in mean games and mean field control problems

Meynard Charles

1) Université côte d'Azur, Charles.meynard@univ-cotedazur.fr

Keywords: Mean field games of controls, Mean field control, Fictitious play, Gradient descent, numerical method.

This abstract is submitted as part of the session "Recents advances in mean field control and mean field games" organised by Pierre Lavigne

In this talk based on an article in preparation [3], we are interested in two learning procedures and their relation to the formation of mean field games (MFGs) equilibria.

The first part is dedicated to the presentation of new results on the fictitious play. Namely, we explore a learning procedure on the controls of players, compared to [1] in which players learn the distribution of the system directly. We prove that this new procedure converges for a class of mean field games of controls (which is to say coefficients depends not only on the distribution of players but also on that of their controls), and possibly with common noise. Besides its interest to explain the formation of mean field games in iterated games, this procedure also suggest a possible alternative approach to learning in MFGs [2, 4].

In the second part, we are concerned with mean field control problems. For such problems, we investigate gradient descent on the space of controls. Under quite general assumptions this gradient descent converges to a mean field game equilibrium, for the associated potential game. This procedure has an interpretation in the formation of MFGs from the optimization problem of a principal agent, moreover it also leads to the presentation of a simple iterative scheme converging to a mean field games equilibrium.

References

[1] Pierre Cardaliaguet and Saeed Hadikhanloo. Learning in mean field games: The fictitious play. ESAIM: Control, Optimisation and Calculus of Variations, 23(2):569–591, April 2017.

[2] Saeed Hadikhanloo, Rida Laraki, Panayotis Mertikopoulos, and Sylvain Sorin. Learning in nonatomic games, part i: Finite action spaces and population games. Journal of Dynamics and Games, 9(4):433–460, 2022.

[3] Charles Meynard. Learning controls in mean field games: the fictitious play and related gradient descent. 2026.

[4] Sarah Perrin, Julien Perolat, Mathieu Laurière, Matthieu Geist, Romuald Elie, and Olivier Pietquin. Fictitious play for mean field games: Continuous time analysis and applications, 2020.

Session 15H: Invited Session:Neural Network approximations for high-dimensional PDEs Problems *Room: Alan Turing (A1.134)*

Numerical bounds for European option prices in uncertainty volatility model.

Jean-François CHASSAGNEUX[1], Simon SANANES[2]

[1]ENSAE-CREST & Institut Polytechnique de Paris, chassagneux@ensae.fr
[2]Université Paris Cité, LPSM & BNP Paribas, simon.sananes@etu.u-paris.fr

Keywords: Dual representation, SPG algorithm, Uncertain Volatility Model

In this talk, we discuss how to produce a numerical interval which contains (or approximately contains) the price of an European option under the Uncertain Volatility Model (UVM). The pricing function, in such models, is solution to a fully non-linear Hamilton-Jacobi-Bellman equation. We first present a numerical method based on stochastic policy gradients for approximating the option price. By construction, this approach provides a lower bound on the true price. For industrial applications, however, this can be unsatisfactory. We therefore construct an upper bound that can also be approximated numerically to provide a computable certificate on the pricing error. The construction is based on PDE residuals and can be interpreted as an a posteriori error estimate.

Session 13A: Geometric Routing *Room: Henri Poincaré (Amphi I)*

Geometric Routing in Geometric Inhomogeneous Random Graphs

Yu-Cheng Chiu, Marc Kaufmann, Kostas Lakis, Ulysse Schaller

1LIX, CNRS, École Polytechnique, IP Paris, France, chiu@lix.polytechnique.fr 2ETH Zürich, Switzerland, marc.kaufmann@inf.ethz.ch, klakis@ethz.ch 3IRIS , University of Stuttgart, Germany, ulysse.schaller@iris.uni-stuttgart.de

Keywords: geometric inhomogeneous random graphs (GIRGs); geometric routing; greedy routing; decentralized algorithms; small-world phenomenon.

We present the first rigorous analysis of decentralized geometric routing in Geometric Inhomogeneous Random Graphs (GIRGs), a model of scale-free networks combining geometry with heterogeneous vertex degrees. Geometric routing is a weight-agnostic variant of the greedy routing protocol. While greedy routing in GIRGs is known to explain the algorithmic small-world phenomenon by finding ultra-short paths of length $\Theta(\log\log n)$ [1], it assumes additional knowledge of vertex weights beyond geometry, an assumption that is often restrictive or unavailable. We investigate whether the underlying geometry alone is sufficient for efficient navigation. We prove that, for power-law weight exponent $\tau \in (2, 3)$ and geometric decay parameter $\alpha > \tau - 1$, geometric routing succeeds with constant probability and finds ultra-short paths of length $\Theta(\log\log n)$, matching the optimal asymptotic guarantees for greedy routing. Our analysis further reveals that, upon success, both protocols follow a similar two-phase trajectory, consisting of a rapid ascent to the heavy vertices, followed by efficient navigation to the target. These results demonstrate that, in the appropriate regime, the network's geometry alone implicitly guides the path to the target through its high-weight core.

References

[1] K. Bringmann, R. Keusch, J. Lengler, Y. Maus, and A. R. Molla, Greedy routing and the algorithmic small-world phenomenon, Journal of Computer and System Sciences, 125:59– 105, 2022.

[2] Y. Chiu, M. Kaufmann, K. Lakis, and U. Schaller, Geometric routing in geometric inhomogeneous random graphs, arXiv:2606.02856, 2026.

Session 8E: Optimisation with chance constraints *Room: Antoine Lavoisier (A1.122)*

ProbLiM: A Julia Toolbox for Linear Problems with Chance Constraints

Yi Chu, Welington de Oliveira, Wim van Ackooij

1Mines Paris-PSL, CMA - Centre de Mathématiques Appliquées, Sophia Antipolis, France 2OSIRIS, EDF Lab Paris-Saclay, France

Keywords: Joint chance constraints, Linear programming, Computational toolbox

We present ProbLiM (Probabilistic Linear Minimization), an open-source Julia toolbox for linear problems involving joint affine chance constraints under elliptically symmetric uncertainty. At its core is a spherical-radial oracle for estimating the joint probability and its gradient under several elliptically symmetric distributions. ProbLiM integrates this oracle with three specialized optimization methods, referred to as the Proximal, Feasible, and Penalty methods, through a unified computational interface. ProbLiM also provides interfaces to a level bundle method and IPOPT, offering alternative approaches for solving the considered problems. We describe the main components of ProbLiM, discuss practical aspects of their implementation, and illustrate its use on several linear problems with joint affine chance constraints. The numerical results highlight the computational efficiency, robustness, and flexibility of the three specialized methods across the considered problem classes.

Session 12E: Convergence rates *Room: Antoine Lavoisier (A1.122)*

On convergence rates of subgradient descent on semialgebraic functions

Evgenii Chzhen, Sholom Schechtman

Evgenii Chzhen: Laboratoire de mathématiques d'Orsay, Université Paris-Saclay, CNRS, 91405 Orsay, France, first.last@cnrs.fr

Sholom Schechtman: SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, 91120 Palaiseau, France, first.last@telecom-sudparis.eu

Keywords: Gradient Descent, Clarke Subgradient, Lipschitz Stratifications, Semialgebraic

We analyze the constant step size subgradient method on nonsmooth, nonconvex functions. We identify geometric assumptions on the objective function under which i) its domain admits a partition (stratification) into smooth manifolds (strata) on which the function is smooth; ii) a global projection formula for Clarke subgradients holds; and iii) quantitative curvature bounds hold on each stratum. Under these conditions, we prove that the iterates of the subgradient method locally shadow a Riemannian gradient descent on nearby strata, which we use to measure stationarity. We introduce a selection rule for the active stratum and develop a mechanism that assembles local descent inequalities across successive strata into explicit convergence rates. These rates are expressed in terms of the number of dimensions present in the stratification, improve as the number of strata decreases, and recover, up to constants, the classical rates in the smooth case. We show that the stated assumptions follow from the existence of Lipschitz stratifications of semialgebraic sets, and are therefore automatically satisfied for semialgebraic functions and, more generally, for functions definable in polynomially bounded o-minimal structures—yielding the first known convergence rates in these settings. As intermediate results of independent interest, we establish tubular neighborhood estimates for Lipschitz stratifications and a global projection formula for Clarke subgradients. Finally, we show that our framework extends to decreasing step size and recovers, via an alternative argument, the recently announced result of Lai and Song on sequential convergence of the subgradient method with step sizes 1/k.

Session 15H: Invited Session:Neural Network approximations for high-dimensional PDEs Problems *Room: Alan Turing (A1.134)*

Deep Hilbert–Galerkin Methods for Infinite-Dimensional PDEs and Optimal Control

Samuel N. Cohen[1], Filippo de Feo[2], Jackson Hebner[1], Justin Sirignano[1]

[1]Mathematical Institute, University of Oxford, {cohens,hebner,justin.sirignano}@maths.ox.ac.uk
[2]Institut für Mathematik, Technische Universität Berlin, defeo@math.tu-berlin.de

Keywords: Infinite-dimensional PDE, Deep Galerkin Method, derivative-informed neural operator, infinite-dimensional control, actor-critic reinforcement learning

We develop deep learning-based approximation methods for fully nonlinear second-order PDEs on separable Hilbert spaces, such as HJB equations for infinite-dimensional control, by parameterizing solutions via Hilbert–Galerkin Neural Operators (HGNOs). We prove the first Universal Approximation Theorems (UATs) which are sufficiently powerful to address these problems, based on novel topologies for Hessian terms and corresponding novel continuity assumptions on the fully nonlinear operator. These topologies are non-sequential and non-metrizable, making the problem delicate. In particular, we prove UATs for functions on Hilbert spaces, together with their Fréchet derivatives up to second order, and for unbounded operators applied to the first derivative, ensuring that HGNOs are able to approximate all the PDE terms. For control problems, we further prove UATs for optimal feedback controls in terms of our approximating value function HGNO. We develop numerical training methods, which we call Deep Hilbert–Galerkin and Hilbert Actor-Critic (reinforcement learning) Methods, for these problems by minimizing the $L^2_{\mu(H)}$-norm of the residual of the PDE on the whole Hilbert space, not just a projected PDE to finite dimensions. This is the first paper to propose such an approach. The models considered arise in many applied sciences, such as functional differential equations in physics and Kolmogorov and HJB PDEs related to controlled PDEs, SPDEs, path-dependent systems, partially observed stochastic systems, and mean-field SDEs. We numerically solve examples of Kolmogorov and HJB PDEs related to the optimal control of deterministic and stochastic heat and Burgers' equations, demonstrating the promise of our deep learning-based approach.

Session 8F: Solving games and Equilibrium problems *Room: Ada Lovelace (A1.128)*

On Stochastic Nonconvex Bilevel Optimization

Daniel Cortild[1], Mathias Staudigl[2], Juan Peypouquet[3], Coralia Cartis[1]

[1]University of Oxford, Oxford, United Kingdom,
[2]University of Economics and Business, Vienna, Austria,
[3]University of Groningen, Groningen, Netherlands,
daniel.cortild@maths.ox.ac.uk, mathias.staudigl@wu.ac.at,
j.g.peypouquet@rug.nl, coralia.cartis@maths.ox.ac.uk

Keywords: Simple Bilevel Programming, Stochastic Optimization, Dynamic Barrier

We investigate stochastic simple bilevel optimization with smooth and possibly nonconvex upperand lower-level objectives. Existing stochastic extensions of dynamic barrier gradient descent (DBGD) [2] either obtain fast convergence under an unverifiable trajectory-dependent "rarevisit" assumption, or remove this assumption at a substantially higher oracle cost [1]. We show that a simple denominator-only regularization of the DBGD multiplier eliminates the need for such an assumption while preserving fast convergence rates [3]. Specifically, our method achieves $(\varepsilon, \varepsilon)$-stationarity in $O(\varepsilon{-}2)$ iterations using $O(\varepsilon{-}4)$ upper-level and $O(\varepsilon{-}7)$ lower-level stochastic gradients, which improves upon the best assumption-free complexities.

References

[1] M. M. Ahmadi, J. Cao, A. Mokhtari, and E. Y. Hamedani. Stochastic Dynamic Barrier Perturbed Gradient Methods for Nonconvex Simple Bilevel Optimization, 2026. arXiv preprint arXiv:2607.10957.

[2] J. Cao, R. Jiang, E. Y. Hamedani, and A. Mokhtari. On the Complexity of Finding Stationary Points in Nonconvex Simple Bilevel Optimization. In Advances in Neural Information Processing Systems, volume 38, 2026.

[3] D. Cortild, M. Staudigl, J. Peypouquet, and C. Cartis. Stochastic Nonconvex Bilevel Optimization: Improved Rates Without Rare-Visit Assumption, 2026. arXiv preprint arXiv:2609.06580.

Session 5K: Decision Focused Learning *Room: Loggia (A1 - first floor)*

Decision-Focused Learning for Dynamic Inventory Replenishment under Production Constraints

Nicolas Corvol[1], Louis Bouvier[1], Axel Parmentier[1]

[1]CERMICS, France, nicolas.corvol@enpc.fr, louis.bouvier@enpc.fr, axel.parmentier@enpc.fr

Keywords: Decision Focused Learning, Combinatorial Optimization, Stochastic Optimization, Revenue Management, Markov Decision Process

We study the dynamic replenishment problem of a single store facing stochastic demand under production quotas. At each period, the store chooses the replenishment quantities of multiple items subject to shared quotas, coupling decisions across the catalog. Demand is endogenous: the available assortment depends on past replenishment decisions and, in turn, influences future demand. The resulting Markov Decision Process has combinatorial state and action spaces and endogenous uncertainty, making dynamic programming intractable. On the other hand, the Sample Average Approximation (SAA) of the corresponding multi-stage stochastic program does not scale with the number of items, limiting its applicability to large catalogs.

We leverage recent advances in contextual stochastic optimization, which learn policies composed of a statistical model followed by a combinatorial optimization layer [1, 2]. The model predicts parameters for the objective of a mixed-integer program, leveraging its structure to learn high-quality and feasible decisions.

We therefore design a combinatorial optimization layer for dynamic replenishment and show that appropriate parameter choices recover many classical inventory policies, including basestock and (s, S) policies and their periodic and capped variants, the capacitated multi-item newsvendor, and constrained assortment optimization under a multinomial logit choice model. We train this policy by extending the primal-dual algorithm for contextual stochastic combinatorial optimization of Bouvier et al. (2026) [3] to the multi-stage setting. Numerical experiments on both academic instances and industrial data from a car manufacturer show that the learned policy achieves performance comparable to SAA while requiring four orders of magnitude less computing time. It also outperforms state-of-the-art inventory management baselines.

References

[1] G. Dalle, L. Baty, L. Bouvier, A. Parmentier, Learning with combinatorial optimization layers: a probabilistic approach, preprint, 2022, arXiv:2207.13513.

[2] L. Baty, K. Jungel, P. S. Klein, A. Parmentier, M. Schiffer, Combinatorial optimizationenriched machine learning to solve the dynamic vehicle routing problem with time windows, Transportation Science, 58(4):708–725, 2024.

[3] L. Bouvier, A. Parmentier, A primal-dual algorithm for contextual stochastic combinatorial optimization, preprint, 2025.

Session 12G: Mid and Long Term Energy Management *Room: Enrico Fermi (A1.133)*

Optimizing Muli-Year Investment Planning for Electricity Distribution Grid Resilience

Axelle Cousinou[1], Sebastien Delbos[1], Mathieu Francois[1]

[1]Enedis, France, axelle-externe.cousinou@enedis.fr, sebastien.delbos@enedis.fr, francois-f.mathieu@enedis.fr

Keywords: network optimization, shortest path algorithm, knapsack problem

Enedis, as the French main DSO, can be severely affected by storms and other extreme weather conditions. In particular, the medium-voltage (MV) network located in wooded areas is prone to be affected by vegetation-related incidents, which may lead to long outages and significant numbers of customers without supply. This paper presents an Operations Research framework developed to transform a national resilience objective into an executable multi-year investment programme. For each MV feeder, all possible electrical networks are generated, including existing electrical assets and feasible underground routes along roads. These networks are combined into weighted graphs. Edge costs represent construction conditions, junctions, substation adaptations, asset removal, and exposure to vegetation-related risks. Shortest-path based algorithms are used to generate network configurations that preserve connectivity while exploring different trade-offs between investment cost and resilience. Each configuration is then assessed in terms of cost and expected reliability. Since feeder-scale configurations cannot be implemented as single construction operations, a graph-based decomposition identifies connected groups of assets that should be removed and uses shortest-path computations to associate the required new sections and substation adaptations. This produces operationally independent projects that preserve the required reliability when other projects are postponed. Each project is characterized by geographical scope, construction actions, investment cost, avoided risk, and reliability benefit. Projects are then selected and scheduled over several years using a budget-constrained knapsack model that maximizes resilience and reliability benefits within annual regional budgets. Project dependencies and shared construction costs are accounted for to produce a consistent multi-year investment plan. Applied to mainland France, this framework generated thousands of projects, with scenarios reducing vegetation-related exposure by more than 70 %, removing approximately 140,000 km of overhead lines, and improving estimated network reliability by 5 10 %. The approach combines network design, Steiner-tree approximation, project decomposition, and multi-period knapsack optimization into an end-to-end decision-support method linking strategic resilience targets to implementable construction programmes.

Session 12A: Invited session: Applications of optimization to power systems and networks *Room: Marie Curie (Auditorium)*

Scheduling of Nuclear Power Plant Outages under Duration Uncertainty

Ezra Daniels, Miguel F. Anjos, Luce Brotcorne, Bárbara Rodrigues, Marius Roland

1University of Southampton, Southampton SO17 1BJ, United Kingdom; 2School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh EH9 3FD, United Kingdom, and GERAD HEC Montréal, Montréal (Québec) Canada H3T 2A7; 3Université de Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL, Lille F-59000, France

Keywords: Nuclear Outage Planning, Two-Stage Stochastic Program, Scheduling of Nuclear units, Energy Systems

We study the Nuclear Outage Planning Problem (NOPP). This problem consists of finding a minimum-cost refuelling outage plan for a fleet of nuclear power plants over a multi-year horizon. Refuelling outages are one of the most complicated phases of a plant's operating cycle [1], requiring thousands of activities to be scheduled within a tight window. Furthermore, these outages are prone to extensions of uncertain duration. An outage that overruns shifts the start of the following production phase, and alters the unit's fuel depletion trajectory. This in turn may raise the cost of meeting demand, delay the outages of other units that share the same resources, or in the worst case leave the fleet without a feasible plan. We propose a two-stage stochastic mixed-integer formulation to decide when to schedule the outages. The first stage selects, for each unit, a path through a transition graph encoding all feasible multi-cycle operational plans [2]. The second stage adapts production dispatch, fuel level and production mode (full power, modulation, or a mandated decreasing power profile) to the realised outage durations in each scenario used to capture uncertainty. Our numerical experiments use the full 58-unit French fleet, with fleet parameters and historical outage durations supplied by EDF and demand taken from RTE's eCO2mix database. The value of the stochastic formulation comes from infeasibility rather than cost. We observe that the generation cost is almost invariant across feasible schedules because the merit order and the reserve margin fix each technology's contribution to demand before any optimisation takes place. What changes instead is whether a feasible schedule exists at all. Feasible schedules of the deterministic instance are found to be infeasible as soon as any duration uncertainty is introduced. This is due to the limit on how many units can be in outage simultaneously being breached. The mechanism used for sampling scenarios determines the frequency and severity of infeasibility. Comonotonic sampling, in which either all outages run as planned, or all overrun by the same percentile, produces infeasibility at the majority of sites but only under extreme realisations. Independent sampling produces infeasibility in the majority of realisations, but only at a few sites. Under both mechanisms the stochastic plan violates the site limits substantially less often than the deterministic one. A limitation of the present numerical experiments is the planning horizon. The computational complexity grows linearly with the scenario count and quadratically with the horizon length. Over a single year most units undergo only one outage, so the propagation of an overrun into a unit's subsequent production cycle is not observed. Capturing this requires a horizon of at least two cycles per unit, and therefore decomposition.

References

[1] S. Lawrence, A. Y. Al Rashdan, D. Mandelli, S. W. St. Germain, N. J. Mapes, C. Wang, and J. J. Cogliati. Tools and methods for optimization of nuclear plant outages. Technical Report INL/RPT-23-74439, Idaho National Laboratory, 2023. URL https://www.osti.gov/biblio/1999587. Prepared for the U.S. Department of Energy, Light Water Reactor Sustainability Program. OSTI ID 1999587.

[2] R. Griset, P. Bendotti, B. Detienne, M. Porcheron, H. Şen, and F. Vanderbeck. Combining Dantzig–Wolfe and Benders decompositions to solve a large-scale nuclear outage planning problem. European Journal of Operational Research, 298(3):1067–1083, 2022.

Session 12H: Market Design and Strategic Behavior *Room: Alan Turing (A1.134)*

Market Equilibrium based Co-Opetitive Infrastructure Sharing

Mandar Datar, Andrea Araldo, Tijani Chahed

1 Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France . firstname.lastname@telecom-sudparis.eu

Keywords: Fisher markets, Co-investment, Variational inequalities, Partition-functionform games

Large-scale infrastructures, such as communication networks, energy systems, and computing platforms, increasingly rely on joint investments by independent stakeholders. Such arrangements enable stakeholders to share deployment and maintenance costs while remaining competitors in the subsequent use of the infrastructure. This creates a coupling between infrastructure investment, resource allocation, and coalition formation: the benefit of forming a coalition depends on both the infrastructure it can jointly finance and the resulting resource allocation. We develop a market-based framework to jointly model these interactions. Stakeholders form co-investment coalitions and pool their budgets to finance infrastructure capacity, which is subsequently allocated among coalition members through a Fisher-market mechanism. Unlike classical Fisher markets with fixed supply, infrastructure capacity is endogenous and determined by the coalition's aggregate investment and production cost. We further allow for allocation externalities, where a stakeholder's utility may depend on both its own allocation and those of other stakeholders, capturing cooperation in infrastructure investment alongside competition in resource utilization. For a fixed coalition structure, we introduce a Partition-Induced Fisher Market Equilibrium (PIFME), which jointly determines infrastructure production, resource allocation, and coalitionspecific prices. We characterize this equilibrium through a variational inequality (VI), where the multipliers associated with capacity constraints recover the equilibrium prices. To study the stability of co-investment structures, we associate each coalition structure with its market equilibrium and formulate a partition-function-form game. Stability is analyzed through the recursive core, which accounts for coalition deviations and the subsequent reorganization of the remaining stakeholders.

References

[1] M. Datar, On stability and learning of competitive equilibrium in generalized Fisher market models: a variational inequality approach, International Conference on Game Theory for Networks, 16(5):42–60, 2025.

[2] Kóczy, László Á, A recursive core for partition function form games, Theory and Decision 63 (1),41–51, 2007.

Session 5K: Decision Focused Learning *Room: Loggia (A1 - first floor)*

Decision-Focused Policies for Microgrid Energy Management

Alexis David[1], Tristan Rigaut[1], Vincent Leclère[2], Adrien Le Franc[2]

[1]Schneider Electric, France, alexis.david1@se.com, tristan.rigaut@se.com
[2,']École des Ponts ParisTech (ENPC), France, vincent.leclere@enpc.fr, adrien.le-franc@laposte.net

Keywords: Microgrid energy management, decision-focused learning, differentiable optimization, stochastic optimization, end-to-end learning

Microgrid energy management is a sequential decision-making problem under uncertainty, where photovoltaic generation and electricity demand are stochastic. Model Predictive Control (MPC) is a widely adopted approach where future net load is first forecasted, typically using machine learning models trained with symmetric losses such as Mean Squared Error (MSE), and then used in a deterministic optimization problem. However, prediction accuracy is not necessarily aligned with operational performance when electricity tariffs induce asymmetric costs. Decision-Focused Learning (DFL) addresses this limitation by training predictive models directly with respect to the downstream optimization objective.

In this work, we compare Predict-then-Optimize (PtO) and direct-policy learning approaches for microgrid control on the EMSx benchmark dataset. We evaluate four learning-based controllers grouped in two families. The PtO family includes conventional MPC trained with a forecasting loss and Decision-Focused Forecasting (DFF), which optimizes forecasting models end-to-end through a differentiable optimization layer. The direct-policy family includes imitation learning of anticipative controller and Learning to Optimize (L2O), which learns battery control actions directly from operating costs through a differentiable projection layer that reconstructs feasibility. We also investigate a recent soft-radial projection mechanism that provides a differentiable closed-form projection onto the battery feasibility set, enabling solver-free training.

Experiments across multiple electricity tariffs show that decision-focused approaches consistently outperform prediction-focused baselines in terms of operating cost. The improvements are particularly significant under tariffs that strongly penalize grid imports. Among the bestperforming methods, L2O combined with soft-radial projection achieves the most competitive operating costs while requiring the lowest training and inference times. These results demonstrate the potential of decision-focused learning as an efficient and scalable alternative for practical microgrid energy management.

Session 12J: Handling dynamics *Room: Emmy Noether (A1.140)*

From Random Games to Deterministic Limits: A Differential-Game Approach to Homogenization

Andrea Davini, Raimundo Saona, Bruno Ziliotto

1Sapienza Università di Roma, Italy andrea.davini@uniroma1.it 2London School of Economics and Political Science, UK r.saona-urmeneta@lse.ac.uk 3Toulouse School of Economics, Université Toulouse Capitole, Institut de Mathématiques de Toulouse, CNRS, France bruno.ziliotto@cnrs.fr

Keywords: stochastic homogenization, differential games, Hamilton–Jacobi equations, random environments

Many optimization and game models evolve in heterogeneous random environments. A natural question is whether, at large space and time scales, microscopic randomness averages out and gives rise to a deterministic effective dynamics. For Hamilton–Jacobi equations, this is the problem of stochastic homogenization, initiated by Lions, Papanicolaou and Varadhan [1]. I will discuss a game-theoretic approach to this question. Via the representation of Hamilton– Jacobi equations as values of zero-sum differential games, homogenization can be viewed as a long-run problem for a game played in a random medium. Our main result [4] establishes homogenization for a class of first-order equations with nonconvex and noncoercive Hamiltonians in environments with finite-range dependence, and provides quantitative convergence estimates. The proof works directly with the game and combines geometric control of trajectories, concentration and approximate subadditivity. It is inspired by the metric/eikonal approach of Armstrong and Cardaliaguet [2] and by the work of Feldman and Souganidis [3]; the probabilistic arguments are close in spirit to techniques from first-passage percolation.

References

[1] P.-L. Lions, G. Papanicolaou, S. R. S. Varadhan, Homogenization of Hamilton–Jacobi equations, unpublished preprint, 1987.

[2] S. N. Armstrong, P. Cardaliaguet, Stochastic homogenization of quasilinear Hamilton– Jacobi equations and geometric motions, J. Eur. Math. Soc. 20(4):797–864, 2018.

[3] W. M. Feldman, P. E. Souganidis, Homogenization and non-homogenization of certain nonconvex Hamilton–Jacobi equations, J. Math. Pures Appl. 108(5):751–782, 2017.

[4] A. Davini, R. Saona, B. Ziliotto, Stochastic homogenization of HJ equations: A differential game approach, Ann. Inst. H. Poincaré C Anal. Non Linéaire 43(4):883–925, 2026.

Session 15C: Energy Systems *Room: André Marie Ampère (Amphi II)*

A Fluid Optimization Approach for Decomposable MDPs with Combinatorial Actions: Application to Offshore Substation Maintenance

Solène Delannoy-Pavy, Axel Parmentier, Vincent Leclère, Manuel Ruiz, Cyrille Vessaire

1Réseau de Transport d'Électricité (RTE), Paris, France firstname.name@rte-france.com 2CERMICS, ENPC, Paris, France firstname.name@enpc.fr

Keywords: Maintenance Optimization, Stochastic Programming

France aims to deploy 45 GW of offshore wind capacity by 2050. The French Transmission System Operator (TSO) is responsible for both building and maintaining the offshore substations that connect these wind farms to the grid, and must pay substantial penalties to producers whenever a substation is partially shut down. Because maintenance costs account for a large share of a substation's lifetime cost, estimating them accurately is essential for investment decisions, such as substation design, that must be made well before the substation enters service. This requires optimizing maintenance operations and evaluating their cost. We model the maintenance of such a substation as a decomposable Markov Decision Process (D-MDP) [1], in which each piece of equipment (transformer, converter, cooling system, etc.) is a subsystem with its own degradation state. While the dynamics decompose across components, maintenance decisions do not: feasible maintenance schedules must satisfy shared operational constraints, such as technician availability and HR rules, while penalties depend on maintenance and degradation across the whole substation. We therefore represent a maintenance schedule as a joint combinatorial action with a non-additive cost. This creates two sources of complexity: the joint state space grows exponentially with the number of components, while the number of feasible joint actions is itself too large to enumerate explicitly. We introduce a fluid linear programming approximation of the MDP that exploits the local structure of the dynamics: each component's evolution depends only on whether it is maintained, while the joint action retains the operational constraints and joint cost. The resulting relaxation avoids the full joint state space but still contains a very large number of action variables, which we handle using column generation. We then derive tighter relaxations by tracking the states and actions of groups of components. Numerical experiments study the trade-off between computational effort and accuracy and demonstrate the performance of the approach. The methodology applies more generally to D-MDPs with decomposable dynamics but joint combinatorial actions and costs.

References

[1] Bertsimas, Dimitris and Miši'c, Velibor V., Decomposable Markov Decision Processes: A Fluid Optimization Approach, Operations Research, 10.1287/opre.2016.1531.

Session 5B: Invited session: Recent Advances in Mean Field Games and Mean Field Control *Room: Henri Poincaré (Amphi I)*

Regularization by Common Noise for Mean Field Games and Mean Field Control

François Delarue, Lukas Wessels

Université Côte d'Azur, France

Keywords: Mean field games, mean field control, common noise, regularization by noise, backward stochastic Hamilton–Jacobi–Bellman equations

In this talk, we demonstrate that an infinite dimensional multiplicative common noise has regularizing properties for mean field games (MFGs) and mean field control (MFC) problems. More precisely, in the case of MFGs, we show without relying on classical monotonicity conditions that the master field U(t,x,m) and its x-derivatives are Lipschitz continuous in the measure argument. As a corollary, in the case of MFC problems, we show that the value function is Lions differentiable which allows us to construct optimal feedback controls. Our proof strategy is based on Girsanov's theorem and a regularity result for backward stochastic Hamilton-Jacobi-Bellman equations. This talk is based on ongoing joint work with F. Delarue.

References F. Delarue, L. Wessels, Regularization by Common Noise for Mean Field Games and Mean Field Control, Work in progress.

Session 8C: Invited session: Mixed Integer Programming and Graphs for Challenging Problems *Room: André Marie Ampère (Amphi II)*

Surrogate-Based Optimization for Adversarial Sonar Placement

Lorenzo Demaria, Axel Kärrolm, Jan Kronqvist

Laboratoire d'Informatique de l'École Polytechnique; KTH Royal Institute of Technology

Keywords: Mixed-Integer Linear Programming, Bilevel optimization, Surrogate optimization, Robust sensor placement

Invited talk in the session "Mixed Integer Programming and Graphs for Challenging Problems".

This work studies adversarial sonar placement for maritime surveillance on discrete hexagonal grids: given a fixed number of sonars, the objective is to maximize the minimum detection cost of an evading submarine that optimally plans a path to a fixed goal. In general, bilevel problems are computationally demanding because they nest a planner inside the decision model and typically lack exact baselines for comparison. However, when the inner problem is a shortest path LP on a graph, the graph incidence matrix is totally unimodular and the follower's problem admits an exact, efficiently solvable LP formulation. This setting provides a provable benchmark for evaluating surrogate approaches. Building on this benchmark, we train compact ReLU surrogates to predict shortest path costs from sensor placement configurations and exploit their piecewise-linearity to embed trained networks exactly into single-level mixed-integer linear programs. Computational experiments show that the resulting surrogate-based MILPs can be solved to obtain robust sensor placements.

We further extend the problem to a setting in which one of the sonars is initially unknown to the evader and its location is revealed when the evader approaches it, allowing the evader to adapt its remaining path. The same surrogate-based pipeline can be used, but the MILP formulation must account for the additional information-dependent decisions. Although an exact formulation can be constructed, solving the extended problem becomes computationally prohibitive at the considered instance sizes. This motivates the development of tractable surrogate-based formulations for the extended adversarial sonar placement problem.

Session 15J: Operations Research for Food and Healthcare *Room: Emmy Noether (A1.140)*

Scheduling for "Ramasse": How to Collect Goods on the Rungis International Market

Natello Descormier, Lucas Berterottière, Pierre Fouilhoux, Massil Hihat, Safia Kedad-Sidhoum

1 Califrais, natello.descormier@califrais.fr 2 Califrais & LPSM, lucas.berterottiere@califrais.fr 3 LIPN, pierre.fouilhoux@lipn.fr 4 Califrais, massil.hihat@califrais.fr 5 CEDRIC, safia.kedadsidhoum@lecnam.net

Keywords: Parallel-machine scheduling, Batch scheduling, Mixed-integer linear programming (MILP), Column generation, Supply chain and logistics

Califrais, the digital operator of the Rungis wholesale market, must nightly collect ("ramasse") ordered goods across market pavilions before delivery trucks depart. We formalize the ramasse as a parallel-machine batch scheduling problem with family setup times, aiming to minimize either maximum delivery completion time or (weighted) average delivery completion time. We establish NP-hardness via reductions to classical scheduling problems, derive structural properties of optimal solutions, lower bounds, and both compact and extended MILP formulations. We then design and compare several solution methods: greedy heuristics with local search, a MILP rounding heuristic with proven approximation guarantees, and a column-generation algorithm on a vehicle-based formulation. Experiments on real Califrais data show column generation achieves solutions of close quality to MILP-based methods at a lower computational cost. We conclude with perspectives on integrating ramasse scheduling with downstream delivery routing.

Session 5B: Invited session: Recent Advances in Mean Field Games and Mean Field Control *Room: Henri Poincaré (Amphi I)*

On the notion of Wasserstein BSDE and its applications

Mao Fabrice Djete[1], Mattia Martini[2]

[1]École polytechnique, CMAP, France, mao-fabrice.djete@polytechnique.edu
[2]École polytechnique, CMAP, France, mattia.martini@polytechnique.edu

Keywords: BSDEs, Mean field control and games, HJB equations on the space of probability measures

Invited session Mean Field Models

Classical Backward Stochastic Differential Equations (BSDEs) theory provides an elegant probabilistic representation of stochastic control problems and their associated PDEs. However, this correspondence breaks down in the mean field (McKean-Vlasov) setting, where the law of the state process enters the dynamics intrinsically, rendering the standard framework inadequate. We discuss a notion of BSDE on the Wasserstein space of probability measures, called a W– BSDE, whose solutions are intrinsically measure–dependent, and its connections with mean field control problems, mean field games, and PDEs on the Wasserstein space. In particular, this provides a probabilistic counterpart to analytic approaches based on the master equation. We finally present a well–posedness result for this class of equations. The key point is that the whole problem can be transferred to a single characteristic equation: whenever the latter admits a solution, both components of the W–BSDE can be reconstructed from it. This talk is based on a joint work with Mao Fabrice Djete.

Session 15K: Multiobjective Optimization and Evolutionary Computation Part II *Room: Loggia (A1 - first floor)*

Understanding the Co-evolutionary Solution Exchange of the MOEA/D: Minimal Communication Gives Maximal Speed-Ups

Benjamin Doerr, Moaad Elmoutassim, Martin S. Krejca, Guy Mushiata Kibulu

1Laboratoire d'Informatique (LIX), CNRS, École Polytechnique, Institut Polytechnique de Paris, France firstname.lastname@polytechnique.edu 2École Polytechnique, Institut Polytechnique de Paris, France firstname.lastname@polytechnique.edu

Keywords: multi-objective optimization, MOEA/D, runtime analysis, speed-up

Besides domination-based algorithms such as the NSGA-II, the decomposition-based MOEA/D is one of the most popular and successful optimization heuristic for multi-objective problems. It decomposes the problem into several, single-objective subproblems. By solving these in a co-evolutionary manner and archiving non-dominated solutions, it approximates or computes the Pareto front. Interestingly, so far no advantages from the co-evolutionary communication of subproblem solutions were proven. In this first mathematical runtime analysis addressing this research question, we investigate how the mutation-based MOEA/D solves the classic OneMinMax benchmark when solutions are exchanged among subproblems in different distances r. We show, for the first time, that the MOEA/D profits from an exchange of solutions, and that already very small communication radii r are sufficient to maximally profit from this feature of the algorithm. More precisely, for all values $r \geq 1$ we prove a tight runtime guarantee of $\Theta(n2)$, where n is the problem size. This is asymptotically faster than the known bound of $O(n2 \log n)$ guarantee for $r = 0$, which we complement with a matching bound of $\Omega(n2 \log n)$.

This work is part of the PGMO-funded project From Mathematical Runtime Analyses to Better Evolutionary Multi-Objective Optimizers (PI: Benjamin Doerr).

Session 15K: Multiobjective Optimization and Evolutionary Computation Part II *Room: Loggia (A1 - first floor)*

Full-Dimensional Many-Objective Benchmarks for Discrete Multi-Objective Optimization

Benjamin Doerr, Martin S. Krejca, Dieks Scholten

1 École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France. firstname.lastname@polytechnique.edu

Keywords: evolutionary multi-objective optimization, many-objective optimization, runtime analysis.

The mathematical runtime analysis of multi-objective evolutionary algorithms has made huge progress in the last years. Yet, most results regard bi-objective problems only. The few existing many-objective benchmarks are mostly obtained by defining classic bi-objective benchmarks on disjoint sets of decision variables. This leads to specific structural properties, among others, that only even numbers m of objectives can occur, and that the Pareto front of such problems is an only m/2 dimensional object, whereas in most practical problems, m-objective problems lead to m −1 dimensional Pareto fronts. In this talk, we propose an alternative way to construct many-objective benchmarks from bi-objective ones. It overcomes the weaknesses of the previous construction but preserves the property that the many-objective benchmark is a natural extension of the bi-objective one. Using techniques similar to [1], we show that the runtime guarantees for the OneMinMax, LeadingOnesTrailingZeros (LOTZ) and OneJumpZeroJumpk benchmarks in the classic construction also hold for our new construction. Because the new construction has the same Pareto front dimension with fewer objectives, the size of the largest incomparable set shrinks, leading to generally tighter bounds. For the LOTZ benchmark, in particular, this tightens the guarantee from $O*((n/m' + 1)2m')$ to $O*((n/m' + 1)m'+1)$, where n is the length of the bit string and m'

the number of blocks. No such bound is known for the classic construction. This work is part of the PGMO-funded project From Mathematical Runtime Analyses to Better Evolutionary Multi-Objective Optimizers (PI: Benjamin Doerr).

References

[1] Simon Wietheger and Benjamin Doerr. Near-tight runtime guarantees for many-objective evolutionary algorithms. Parallel Problem Solving from Nature, PPSN 2024, Part IV, pages 153-168, Springer, 2024.

Session 8I: Decision making under uncertainty *Room: Sophie Germain (A1.139)*

Graph powers for robust optimization

Alexandre Dupont-Bouillard, Leo Liberti, Antonio Mucherino

1Université de Rennes IRISA CNRS UMR 6074 , France, dupont-bouillard@lipn.fr 2Institut polytechnique de Paris LIX CNRS UMR 7161, Palaiseau,

Keywords: Graph powers, robust optimization, complexity theory, integer programming

This presentation proposes to study powers of graphs as robust structures. The first part aims at tackling complexity-theoretic questions related to powers of subgraphs and the second focuses on integer programming for solving related problems while keeping an eye on power of Hamiltonian cycles as a case study.

The k-th power Gk of a graph G = (V, E) is the graph obtained from G by adding edges between pairs of vertices at distance at most k in G. Seeking a power of a given subgraph might provide a more robust solution than searching for the subgraph itself: we first establish the robustness guarantee provided by the k-th power of a Hamiltonian cycle in terms of edge deletion while deleting any edge in a Hamiltonian cycle yields a non-Hamiltonian graph.

Given an existence problem Π such as "Does the graph G contain a subgraph satisfying property P" for some fixed property P, we consider its empowered variant ΠR: "Does the graph G contain a subgraph H such that $\sqrt[k]{H}$ exists and satisfies property P". While computing the power of a graph is doable in polynomial time, checking whether $\sqrt[k]{G}$ exists is NP-hard [3, 2]. Therefore, for some property P and encoding of H might not be a polynomially checkable certificate [2] and therefore might be harder than checking P itself. However, checking whether $\sqrt[k]{H}$ satisfies P and whether G contains H is as hard as checking P (when P is nontrivial). We formalize this idea and show that ΠR is no harder than Π in the polynomial hierarchy, making it compatible with other robustness criteria. Based on the theory of SSP-reduction [1] we also establish Σ_2^p-completeness of the problem of finding a min-max-interval-regret-k-th-power of Hamiltonian cycle.

Finally, we present several integer linear programs for solving the problem of finding a minimum weight k-th power of a Hamiltonian cycle. Moreover, we propose a dedicated formulation for the special case k = 2 and show experimentally that it outperforms the generic formulations on TSPLIB instances. However, it is unclear how to generalize it to other values of k. We show that its naïve generalization already fails for k = 3.

References

[1] Grüne, C., and Wulf, L. Completeness in the polynomial hierarchy for many natural problems in bilevel and robust optimization. In Integer Programming and Combinatorial Optimization (Cham, 2025), N. Megow and A. Basu, Eds., Springer Nature Switzerland, pp. 256–269.

[2] L. C. Lau and D. G. Corneil, Recognizing powers of proper interval, split, and chordal graphs. SIAM Journal on Discrete Mathematics 18(1) (2004), 83–102.

[3] R. Motwani and M. Sudan, Computing roots of graphs is hard. Discrete Applied Mathematics 54(1) (1994), 81–88.

Session 12F: Routing & Mobility *Room: Ada Lovelace (A1.128)*

Pricing for Real-Time Ride-Sharing with Random Demand and Supply

Rose Sossou Edou, Virginie Boutueil, Frédéric Meunier

rose.sossou-edou@enpc.fr,virginie.boutueil@enpc.fr, frederic.meunier@enpc.fr
École nationale des ponts et chaussées, France

Keywords: dynamic pricing; ride-sharing; stochastic demand; stochastic supply; simulation; intermediate transport.

Ride-sharing (also known as carpooling) is an emblematic form of intermediate transport. Achieving a satisfactory level of service raises several challenges, particularly that of incentivizing drivers to share their rides. Local initiatives have emerged to support the development of this form of intermediate transport. This work aims to provide methodological tools for public operators seeking to implement monetary incentives for drivers. It originates from discussions with Ecov, a French mobility operator committed to solidarity-based transport. This problem has received limited attention in the operations research literature. Existing studies address casual carpooling [1] and dynamic pricing in large stochastic ride-sharing markets through asymptotic models [2]. Our setting instead considers pre-planned driver trips, incentivedependent acceptance of waiting passengers, and a limited operator budget. As a first step, we propose a streamlined model consisting of a single route with a unique origin–destination pair. Passengers arrive according to an inhomogeneous Poisson process and are willing to wait for a seat for a limited amount of time. Drivers planning to travel along the same route also arrive according to an inhomogeneous Poisson process. Both arrival processes are assumed to be fixed and exogenous. Upon arrival, a driver accepts to share the ride with a waiting passenger with a probability that increases with the incentive offered by the operator at that time. We propose several heuristics for determining the incentive levels, mainly based on approximate deterministic formulations of the problem. The simplest formulation, although potentially less efficient, leads to interesting two-dimensional non-convex optimization problems that can nevertheless be solved using efficient algorithms. A discrete-event simulator is currently being developed to compare and evaluate the different pricing strategies. The objective is to determine when dynamic pricing can improve the number of passengers served and when its effectiveness is limited by stochastic demand and supply, driver acceptance behavior, and budget constraints.

References

[1] S. Cui, K. Li, L. Yang, and J. Wang, Slugging: Casual carpooling for urban transit, Manufacturing & Service Operations Management, vol. 24, no. 5, pp. 2516–2534, 2022.

[2] J. M. Cashore, P. I. Frazier, and É. Tardos, Dynamic pricing provides robust equilibria in stochastic ridesharing networks, Mathematics of Operations Research, vol. 49, no. 3, pp. 1647–1677, 2024.

Session 15B: Learning & Neural Networks *Room: Henri Poincaré (Amphi I)*

On the Explainability of Max-Plus Neural Networks

Ikhlas Enaieh, Olivier Fercoq, Angel Garcia Pedrero

1LTCI, Telecom Paris, Institut Polytechnique de Paris, France, ikhlas.enaieh@telecom-paris.fr, olivier.fercoq@telecom-paris.fr 2DATSI, Universidad Politecnica de Madrid, Madrid, Spain, angel.garcia@ctb.upm.es

Keywords: Max-plus neural networks, explainable artificial intelligence, subgradient optimization, k-medoids

We investigate the explainability properties of the recently proposed linear-min-max neural networks [1]. At initialization, they can be interpreted as k-medoids with the infinity norm as a distance [2, 3]. Then, they are trained using subgradient descent to better fit the data [4]. The model has been shown to be a universal approximator [1]. Yet, we can trace the decision process because a single most activated neuron is responsible for the value of the output. Using this property, we designed a pixel fragility measure [2] that determines whether changes to a single pixel may be responsible for a change in the classification output. Experiments on the PneumoniaMNIST dataset [5] show that this explanation for the output of the neural network compares favorably to SHAP [6] and Integrated Gradients [7].

References

[1] Y. Luo and S. Fan, Min-Max-Plus Neural Networks, arXiv:2102.06358, 2021.

[2] I. Enaieh, O. Fercoq, and A. Garcia Pedrero, On the explainability of max-plus neural networks, IEEE International Symposium on Computer-Based Medical Systems (CBMS), 2026, hal-05601301.

[3] L. Kaufman and P. J. Rousseeuw, Clustering by means of medoids, in Statistical Data Analysis Based on the L1-Norm and Related Methods, pp. 405–416, 1987.

[4] I. Enaieh and O. Fercoq, Exploiting Subgradient Sparsity in Max-Plus Neural Networks, HAL preprint, 2026, hal-05502128.

[5] J. Yang et al., MedMNIST v2: A large-scale lightweight benchmark for 2D and 3D biomedical image classification, Scientific Data, 10:41, 2023.

[6] S. M. Lundberg and S.-I. Lee, A Unified Approach to Interpreting Model Predictions, Advances in Neural Information Processing Systems, 30, 2017.

[7] M. Sundararajan, A. Taly, and Q. Yan, Axiomatic Attribution for Deep Networks, in Proceedings of the 34th International Conference on Machine Learning, pp. 3319–3328, 2017.

Session 5A: Invited session: Nonconvex landscapes and algorithms in data science *Room: Marie Curie (Auditorium)*

Benign Landscape Of Nonconvex Low-dimensional Relaxation Of Orthogonal Synchronization With Group-symmetric Dual Certificates

Faniriana Rakoto Endor[1], Irène Waldspurger[2]

[1]Université Paris-Dauphine and CNRS, France, rakotoendor@ceremade.dauphine.fr
[2]École Normale Supérieure and ENS, France, Irene.Waldspurger@ens.fr

Keywords: Non-convex optimization, Benign landscape, Synchronization

Orthogonal synchronization is the problem of estimating orthogonal matrices from relative pairwise measurements. This problem admits a convex semidefinite program as a relaxation, which is solvable in theory but not scalable in practice. The Burer-Monteiro factorization of this convex program gives a more scalable alternative at the cost of losing convexity. In practice however, local algorithms such as gradient descent find global optima of the non-convex problem, which suggests that, in some case, the landscape is benign : all local minima are global minima. Recent years has seen some results explaining why the landscape is benign, see [1, 2] for orthogonal synchronization and [3, 4] for the particular case of synchronization on spheres.

The goal of this presentation is to understand why that is the case. More precisely, we transform the problem of "explaining why there are no spurious local minima" into an optimization problem (P), the value of which we aim to lower bound. We then discuss how to construct a family of manageable dual certificates for the dual of (P) yielding benign landscape results.

References

[1] S. Ling, Local geometry determines global landscape in low-rank factorization for synchronization, Found. Comput. Math. 26(3) (2026), 1553–1585, arXiv:2311.18670.

[2] S. Ling, Improved global landscape guarantees for low-rank factorization in synchronization, preprint (2026), arXiv:2601.20292.

[3] A. D. McRae, Benign landscapes for synchronization on spheres via normalized Laplacian matrices, preprint (2025), arXiv:2503.18801.

[4] F. Rakoto Endor, I. Waldspurger, Benign landscape for Burer–Monteiro factorizations of MaxCut-type semidefinite programs, preprint (2024), arXiv:2411.03103.

Session 12C: Invited session: Geometric Control and Optimization *Room: André Marie Ampère (Amphi II)*

Optimal control under uncertainty for relative guidance during orbital rendezvous

Laura Bernardo Estévez, Prince Edorh, Bruno Hérissé, Frédéric Jean

1Unité de Mathématiques Appliquées, ENSTA, Institut Polytechnique de Paris, Palaiseau 91120, France, laura.bernardo@ensta.fr, frederic.jean@ensta.fr 2DTIS, ONERA, Université Paris Saclay, Palaiseau 91123, France, prince.edorh@onera.fr, bruno.herisse@onera.fr

Keywords: Optimal control, orbital rendezvous, state constraints

Orbital rendezvous [1] requires efficient maneuvers while ensuring collision avoidance throughout the trajectory [2]. We consider this problem in the context of ONERA's FlyLab mission, involving two nanosatellites in low Earth orbit with different propulsion systems and limited communication capabilities. These constraints lead to an open-loop relative guidance problem, which we formulate as an optimal control problem aiming to minimize the fuel consumption of the chaser satellite and to reach a final configuration while satisfying several operational constraints as well as a minimum safety distance constraint [3]. In this talk, we present the modeling of the problem and introduce a simplified state-constrained formulation that allows us to analyze the structure and derive some qualitative properties using the Pontryagin Maximum Principle (PMP) [4].

References

[1] David A Vallado. Fundamentals of astrodynamics and applications, volume 12. Springer Science & Business Media, 2001.

[2] Simone D'Amico. Autonomous formation flying in low earth orbit. PhD thesis, TU Delft, 2010.

[3] Jihe Wang, Chenglong Xu, Chengxi Zhang, and Jinxiu Zhang. Passive safety-constrained impulsive maneuvers for formation reconfiguration: An analytic solution. Aerospace Science and Technology, 158:109901, 2025.

[4] Richard B Vinter and RB Vinter. Optimal control, volume 2. Springer, 2010.

Session 5C: IS: Modelling supply chains in Mathematical Programming: flow and economics *Room: André Marie Ampère (Amphi II)*

An AMM for Energy Sharing: From Theory to Verifiable Settlement

Michele Fabi[1], Baptiste Leloup[2]

[1]Télécom Paris and CREST, IP Paris, France, michele.fabi@telecom-paris.fr
[2]ENSAE Paris, France, baptiste.leloup@ensae.fr

Keywords: Local energy market, automated market maker, mean-field game, blockchain, zeroknowledge proof

Distributed generation turns households into prosumers, who sell their surplus to the grid at the feed-in tariff and buy their deficit at the retail rate. A local energy community exists to recover the gap between the two, since any exchange between neighbours at a price inside the band benefits both sides. Two questions follow. How to set the price, and how to execute the mechanism while extending as little trust as possible to any party, including the market operator (an entity the French framework makes mandatory). This talk presents the theoretical framework that answers the first question and the proposed infrastructure that answers the second.

Theoretical framework. We develop an axiomatic theory of automated market makers (AMMs) for energy sharing [1]. Each fifteen-minute session matches the aggregate supply and demand of the community at a price inside the tariff band and settles the residual with the grid. The axioms (anonymous execution, budget balance, coalition-proofness) single out linear payment functions whose prices decrease in the supply-to-demand ratio. Prosumers optimise consumption, storage and trade in a dynamic game, which the AMM turns into a potential game: its equilibrium coincides with the social planner's optimum.

Proposed infrastructure. Blockchain and smart contracts are the natural executors to reduce that trust. But the mechanism needs each household's quarter-hourly net position, which reveals occupancy, and writing it to a permanent ledger is hardly compatible with the GDPR. We present a settlement layer that runs the AMM on a permissioned Hyperledger Besu network under QBFT consensus while publishing no individual position. The operator publishes only community aggregates during the day and one cryptographic commitment per member at midnight. Once a day, a zero-knowledge proof establishes that the hidden settlement is the one the public prices imply, that no balance goes negative, and that money is conserved. On ten physically simulated dwellings with real day-ahead prices, the community gains 8.1% relative to the grid-only regime and a day is proven in about five seconds.

References

[1] M. Fabi, V. Nadkarni, L. Leone, M. X. V. Ferreira, Automated market making for energy sharing, working paper, December 2025, arXiv:2512.24432.

Session 5J: Invited session : Quantum Applications in Operations Research part 1 *Room: Emmy Noether (A1.140)*

A mixer for minimization of function with constraint

Gerard Fleury, Philippe Lacomme

Université Clermont-Auvergne, Clermont-Auvergne-INP, LIMOS (UMR CNRS 6158), Aubière, France

Stabilizer theory defines another novel perspective on mixers. We propose to introduce the main definitions and provide an application framework for this theory to design mixers describing potential solutions in the context of constrained function minimization. Building upon the work of [1][2], we formalize an algorithmic description of the theory to enable its practical application to optimization problems. The key advantage is to construct a superposition restricted solely to quantum states satisfying the constraints, thereby achieving an efficient modeling. We propose an iterative constraint- traversal method that iteratively generates the mixer space at each step. Numerical examples are provided to model systems with 5 constraints.

References

[1] Fuchs F.G. and Ruben Pariente Bassa. LX-mixeurs for QAOA: Optimal mixeurs restricted to subspaces and the stabilizer formalism. arXiv:2306.17083v8. 2024.

[2] Fuchs F.G., K.O. Lye, H. M. Nilsen, A.J. Stasik and G. Sartor. Constraint Preserving Mixeurs for the Quantum Approximate Optimization Algorithm. Algorithm. Algorithms, 15: 202, 2022.

Session 15K: Multiobjective Optimization and Evolutionary Computation Part II *Room: Loggia (A1 - first floor)*

Specification of the GDEMO-II Variant for Box Archiving

Jules Fraizier, Benjamin Doerr, Martin S. Krejca

1Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France, jules.fraizier@telecom-paris.fr 2Laboratoire d'Informatique (LIX), CNRS, École Polytechnique, Institut Polytechnique de Paris, France, firstname.lastname@polytechnique.edu

Keywords: Multi-objective optimization, evolutionary algorithms, box archiving, GDEMO, runtime analysis

Multi-objective evolutionary algorithms (MOEAs) are among the most widely and successfully applied optimizers for multi-objective problems. However, because the true Pareto front can be prohibitively large, it is necessary to bound the population size while still computing a good approximation of the optimal trade-offs. MOEAs frequently utilize grid-based box archiving, parameterized by a grid size δ, to compute an additive ϵ-approximation of the Pareto front [LLY24]. Yet, standard archiving policies that strictly retain one individual per box, such as the Global Diversity Evolutionary Multi-Objective Optimizer (GDEMO), suffer from a structural vulnerability. Driven by the bias of standard mutation toward central solutions, unguided intra-box replacement fails to preserve extreme solutions, preventing the algorithm from progressing into new boxes. This traps algorithms in exponential runtime bottlenecks on boundary-critical pseudo-Boolean problems [HN08]. To this end, we propose GDEMO-II, a refined evolutionary archiver. It introduces a global filter that rejects a new offspring only if its targeted box is strictly dominated by the combined population, and it abolishes the single-occupant restriction to explicitly retain the local boundary extremes for each objective. Via mathematical runtime analysis, we prove an upper bound on the population size that takes into account the range of the objective functions and the grid parameter δ. For bi-objective problems where objective values share the same asymptotic bound n, this guarantees a linear memory complexity of $\mu = O(n/\delta)$ Furthermore, we prove that GDEMO-II completely resolves the exponential delay of the original GDEMO on the OneMinMax benchmark, computing a tight δ-approximation in expected time O(n2 log n/δ). Finally, it efficiently discovers the Pareto front of the LOTZ benchmark in expected time O(n3/δ).

References

[HN08] Christian Horoba and Frank Neumann. Benefits and drawbacks for the use of epsilon-dominance in evolutionary multi-objective optimization. In Genetic and Evolutionary Computation Conference, GECCO 2008, Proceedings, pages 641–648. ACM, 2008.

[LLY24] Miqing Li, Manuel López-Ibáñez, and Xin Yao. Multi-objective archiving. IEEE Trans. Evol. Comput., 28(3):696–717, 2024.

Session 5H: The moment hierachy and SDP *Room: Alan Turing (A1.134)*

Some curiosities about the convergence of the Moment-SOS hierarchy

Adrien Le Franc

1CERMICS, École nationale des ponts et chaussées, IP Paris, France, first-name.last-name@enpc.fr

Keywords: Polynomial Optimization, Semi-definite Programming, Moment-SOS Hierarchy

Polynomial optimization encompasses problems for which the objective function and constraints are encoded by multivariate polynomials. In the general case, such problems may have multiple local minima and can be numerically challenging. Given a specific polynomial optimization problem (POP), the moment-SOS hierarchy introduces a sequence of semi-definite programs whose values converge to the infimum of the POP [1]. This approach has shown promising numerical performances, and technical conditions to enforce convergence in finitely many steps of hierarchy were given in [2]. Yet, even when finite convergence holds, there still lacks a systematic understanding of the amount of steps required for the hierarchy to converge. In this talk, we introduce a parametrized quadratic univariate POP for which the momentSOS hierarchy always converges in finitely many iterations. Interestingly, the number of steps required to achieve convergence can be arbitrarily large, as the parameter tends to a limit value [3]. In the second part of the talk, we also discuss some examples of converging sparse momentSOS hierarchies without the standard running intersection property.

References

[1] J.-B. Lasserre. Global optimization with polynomials and the problem of moments. SIAM Journal on Optimization, 2001.

[2] J. Nie. Optimality conditions and finite convergence of Lasserre's hierarchy. Mathematical Programming, 2014.

[3] D. Henrion, A. Le Franc, V. Magron. Slow convergence of the moment-SOS hierarchy for an elementary polynomial optimization problem. SIAM Journal on Applied Algebra and Geometry, 2025.

Session 12D: Optimal Transport & Distributional Robustness *Room: Lise Meitner (A1.116)*

Exact Reformulation for Nonconvex Wasserstein DRO: Application to Day-Ahead Renewable Nomination

Charles Garrisi, Welington de Oliveira, François Pacaud

Mines Paris PSL

Keywords: distributionally robust optimization, Wasserstein distance, electricity markets

Day-ahead renewable nomination requires committing energy before both generation and settlement prices are known. Wasserstein distributionally robust optimization (DRO) provides a natural framework for hedging against uncertainty, especially when historical data are scarce. However, when both the power generation and settlement prices are unknown, the objective function becomes nonconvex in the uncertainty. This makes the standard tractable convex reformulation of Wasserstein DRO inapplicable directly. A common technique to restore tractability consists in linearizing the loss at the cost of excessive conservatism, which we can address via an appropriate McCormick relaxation. An alternative approach is to directly solve the original nonconvex problem. By exploiting the low-dimensional structure induced by a box support and a weighted$\ell_1$transportation cost, we prove that each sample-wise separation problem attains its optimum on finitely many uncertainty states, yielding an exact finite reformulation. Synthetic experiments show that distributional robustness matters when limited or highly variable data misestimate marginal imbalance incentives. Experiments on actual French market data confirm that distributionally robust methods can tame exposure to extreme imbalance episodes, at the cost of lower average profit.

Session 15A: Optimal Control, Measures and PDE constrained optimization *Room: Marie Curie (Auditorium)*

Adjoint-based Optimality Conditions for a Probability Maximization Problem

Caroline Geiersbach, Wim van Ackooij, Hasnaa Zidani

1University of Klagenfurt, Austria, caroline.geiersbach@aau.at 2Électricité de France (EDF R&D), France, wim.van-ackooij@edf.fr 3LMI, INSA Rouen Normandie, Rouen, France, hasnaa.zidani@insa-rouen.fr

Keywords: Optimal control under uncertainty, Pontryagin-type optimality conditions, probability functions, lookback problem

This talk discusses new results building on the publications [1] and [2], where probability functions have been analyzed in the context of optimal control problems under uncertainty. We consider a linear parabolic partial differential equation in which the control appears as a distributed source term and uncertainty acts in a structured way on the right-hand side. The objective is find a control such that some outcome of the random state occurs with maximal probability. The outcome can be enforced pointwise, for instance when requiring the state to remain bounded at every point in time and space. This type of model offers a compromise between the robust "lookback" problem and the problem of maximizing on average. Probability functions are attractive because of their interpretability, but are generally nonsmooth and nonconvex. In the case of Gaussian-like noise, the properties of the underlying control-to-state operator allow us to obtain a fully explicit formula for the Clarke subdifferential of the probability function. The formula depends on certain solutions to an adjoint equation with measures appearing in source terms. We also consider optimality conditions depending on solutions to a related adjoint equation in the case of the lookback problem. These formulae are used to compute solutions in a numerical example.

References

[1] C. Geiersbach and R. Henrion, Optimality conditions in control problems with random state constraints in probabilistic or almost sure form, Mathematics of Operations Research, 2024.

[2] W. van Ackooij, R. Henrion, and H. Zidani, Pontryagin's principle for some probabilistic control problems, Applied Mathematics & Optimization, 90(1), 2024.

Session 15F: Combinatorial Optimization & Applications *Room: Ada Lovelace (A1.128)*

Combinatorial Optimization for Energy-Efficient In-Network Learning

Garance Gérard, Mustapha Bouhtou, Safia Kedad-Sidhoum, Amélie Lambert, Nancy Perrot

1Orange Research, France, garance.gerard@orange.com, mustapha.bouhtou@orange.com, nancy.perrot@orange.com
2Cedric-Cnam, France, garance.gerard@lecnam.net, safia.kedad_sidhoum@cnam.fr, amelie.lambert@cnam.fr

Keywords: Combinatorial Optimization, Valid Inequalities, In-Network Learning

Machine learning applications may involve data distributed across multiple sources rather than collected at a central location. In-network learning addresses this setting by deploying machine learning and data processing functions within the network, closer to the data sources. Deciding where functions are placed, how they are chained from sources to aggregators, and how the resulting flows are routed leads to a combinatorial optimization problem in which chaining, assignment, and routing decisions are coupled. These decisions directly affect network energy consumption, motivating the design of energy-efficient function placement and routing strategies. Building on the Data Pipeline System framework of [1], we formulate this joint chaining, assignment, and routing problem as a Quadratically Constrained Quadratic Program that minimizes energy consumption, subject to node and link capacity constraints and end-to-end latency requirements. We prove NP-hardness by reduction from the Uncapacitated Facility Location Problem. We derive families of valid inequalities and strengthening constraints to tighten the formulation. We also investigate alternative formulations and decomposition approaches to improve scalability on larger instances. Computational experiments on SNDlib topologies, with varying numbers of data sources, quantify the impact of latency constraints on optimal energy consumption and assess the computational performance of the proposed approaches. This work extends the model presented at INOC 2026 [2].

References

[1] P. Ntumba, N.-E.-H. Yellas, S. Bin-Ruba, F. Ben Abdesslem, and S. Secci, Data Pipeline System Designs for In-network Learning, 2024 20th International Conference on Network and Service Management (CNSM), 2024, pp. 1–9.

[2] G. Gérard, P. Ntumba, S. Kedad-Sidhoum, A. Lambert, and N. Perrot, Energy-Efficient Function Chaining and Assignment for In-Network Learning, INOC 2026, Liège, Belgium, HAL:05556870.

Session 8C: Invited session: Mixed Integer Programming and Graphs for Challenging Problems *Room: André Marie Ampère (Amphi II)*

Graph Partitioning and ADMM for Optimal Pump Scheduling in Drinking Water Distribution Networks

Carlo A. M. Gervasi, Stefan Ropke, Claudia D'Ambrosio, Bissan Ghaddar

1Technical University of Denmark, Kongens Lyngby, Denmark, EMAIL@dtu.dk, ropke@dtu.dk 2LIX, CNRS & Ecole Polytechnique, Institut Polytechnique de Paris, France, dambrosio@lix.polytechnique.fr 3IE University, Madrid, Spain, bghaddar@faculty.ie.edu

Keywords: decomposition methods, ADMM, mixed-integer nonlinear programming, graph partitioning, water networks

Problem. Optimal pump scheduling in Drinking Water Distribution Networks (DWDNs) seeks the pump on/off profile over a discrete horizon minimising energy cost subject to demand, tank level and pressure requirements. The formulation is a Mixed-Integer Nonlinear Program: head losses follow the nonconvex Hazen Williams law, pump curves are nonlinear, and binary scheduling variables are coupled across time by tank dynamics. Monolithic solution is impractical beyond small benchmarks; the literature has responded either with simplified hydraulics, which weakens feasibility guarantees, or with metaheuristics, which give no optimality information. Approach. We study a decomposition-based framework that keeps the full hydraulic model and distributes the computational effort instead of simplifying the physics. Graph partitioning is applied to the network topology to produce subnetworks joined by a small set of boundary nodes; the intuition, which we test rather than assume, is that hydraulic coupling concentrates on few links, so a good cut yields balanced and weakly coupled subproblems. The Alternating Direction Method of Multipliers (ADMM) then acts as the coordination layer: each subnetwork solves its own MINLP and consistency of heads and flows at the boundary nodes is enforced through the dual updates, with binary pump variables handled inside the loop rather than relaxed away. Caveats. ADMM has no convergence guarantee on a nonconvex mixed-integer problem, so the iterates are used as a heuristic: the output is post-processed into a feasible schedule and assessed against reference solutions, and what we report is a measured gap, not a certified bound. Performance is also sensitive to the partition number of parts, balance criterion, edge weighting and to the penalty parameter ρ, with no single choice dominating across networks. Characterising which topological features make a DWDN amenable to this decomposition is, in our view, the open question rather than a settled step. Experiments and discussion. The framework is evaluated on benchmark networks of increasing size and topological complexity, against the monolithic MINLP where tractable and against a non-partitioned baseline, reporting solution cost, constraint violation after recovery, and wall-clock time, together with a sensitivity study over the number of partitions and ρ. The talk focuses on where the decomposition helps, where it does not, and what the observed failure modes reveal about the interaction between network structure and distributed optimisation. Directions under investigation include partition selection informed by hydraulic rather than purely topological edge weights, and warm-starting the subproblems across a rolling horizon.

Session 15D: Optimisation challenges in modern mobility *Room: Lise Meitner (A1.116)*

Online Combinatorial Learning with Causal Rescheduling for EV Charging

Eymeric Giabicani[1], Raphaël Féraud[4], Patrick Maillé[2], Guy Camilleri[3], Anne Blavette[1]

[1]Univ Rennes, ENS Rennes, CNRS, IETR, France
eymeric.giabicani@ens-rennes.fr, anne.blavette@ens-rennes.fr
[2]IMT Atlantique, France
[3]IRIT, Toulouse, France
[4]Orange Innovation, France

Keywords: Online Learning · Combinatorial Bandits · Causal Rescheduling · EV Charging

Online learning usually assumes that the action selected by a learner is also the action that is executed. In safety-critical systems, protection mechanisms may instead invalidate part of a decision during execution, so that the realized action, loss, and feedback differ from the initial plan. We study this execution mismatch in decentralized electric-vehicle (EV) charging. Each EV selects a combinatorial charging schedule within a stochastic availability window, while transformer-protection events may interrupt charging and force causal rescheduling to future feasible slots. We introduce Causal Rescheduling Regret (CR-Regret), which compares the executed trajectory with the best fixed causal deployment policy in hindsight. Our main result is a structural transfer theorem from adversarial combinatorial semi-bandits to this execution-aware regret notion. If $r_t \in \{0,1\}^d$ denotes the interruption mask and $F_T = \sum_{t=1}^{T} \|r_t\|_1$, then for any admissible online initial-selection algorithm,

$$\text{CR– Regret}(T) \leq \text{Regret}(T) + \mathbb{E}[F_T].$$

Hence, guarantees for the underlying combinatorial bandit learner transfer to the rescheduling problem with only an additive execution-mismatch term, independently of the state-space size of the induced recourse MDP. Combining this framework with SleepingCatBandit for stochastic decision sets and adversarial semibandit losses [1] yields FTPL-Repair, with

$\widetilde{O}((mdT)^{2/3}) + O(F_T)$

CR-Regret, which is sublinear whenever $F_T = o(T)$. We evaluate the method in controlled experiments and in a fleet-scale study using four years of photovoltaic-generation and electricity-consumption data. With 480 EVs, causal rescheduling maintains 100% demand satisfaction. A safety-oriented FTPL-Repair variant reduces congestion by more than 99% relative to immediate charging, with fewer than two disconnections per EV, while a lower-price variant remains competitive with a centralized forecast-based optimizer. These results illustrate how executiontime safety interventions can be incorporated into adversarial online optimization while preserving regret guarantees.

References

[1] G. Neu and M. Valko, Online combinatorial optimization with stochastic decision sets and adversarial losses, Advances in Neural Information Processing Systems, 2014.

Session 5F: IS : Recent Advances in Systematic and Computer-Aided Methods in Optimization I *Room: Ada Lovelace (A1.128)*

Limited-Memory First-Order Methods and H -Duality

Pierre Vernimmen[1], François Glineur[1]

[1]UCLouvain, Belgium, pierre.vernimmen@uclouvain.be, françois.glineur@uclouvain.be

Keywords: Numerical Algorithm Design, Limited-Memory Algorithm, H-duality

A fixed-coefficient first-order method performing N gradient evaluations can be represented as

$$x_{k+1} = x_k - \frac{1}{L}\sum_{i=0}^{k} H_{k,i}\nabla f(x_i), \qquad H \in \mathbb{R}^{N\times N},$$

where H is a lower-triangular matrix. Recent works design optimal methods by numerically optimizing coefficients in H [1, 2]. In this representation, computing an iterate may require storing all past gradients. However, many methods, including standard accelerated methods, admit a much more compact recurrence that stores only a few summaries of the past. In this work we focus on memory-r methods, where the r stored vectors are collected in a state matrix $S_k \in \mathbb{R}^{d\times r}$ and updated linearly:

$$x_{k+1} = x_k - \frac{1}{L}(\nabla f(x_k)d_k + S_k c_k), \qquad S_0 = 0, \quad S_{k+1} = S_k A_k + \nabla f(x_k)b_k,$$

with $A_k \in \mathbb{R}^{r\times r}$, $b_k \in \mathbb{R}^{1\times r}$, $c_k \in \mathbb{R}^{r\times 1}$, and $d_k \in \mathbb{R}$ for all k. We show that the minimum memory required to implement the method described by H is $r = \max_k \operatorname{rank} H_{k:N-1,\,0:k-1}$, and we give an algorithm to recover (A_k, b_k, c_k, d_k) from H.

The H-dual of a method, obtained by anti-transposing matrix H, was shown in some situations to transfer convergence guarantees between different performance metrics [3, 4]. We show that taking the H-dual of a memory-r method gives another memory-r method described by

$$(A_k, b_k, c_k, d_k) \overset{\text{H-dual}}{\longleftrightarrow} (\widetilde{A}_k, \tilde{b}_k, \tilde{c}_k, \tilde{d}_k) = (A_{N-1-k}^\top, c_{N-1-k}^\top, b_{N-1-k}^\top, d_{N-1-k}) \quad \text{for all } k.$$

Shu and Wang [4] recently considered the design of an optimized method for L-smooth μ-strongly convex functions, targeting the terminal distance relative to the initial gradient norm. For $N = 5$ and $q = \mu/L = 1/11$, we show that their numerically optimized method is memory-1, up to rounding errors. Reconstructing its compact recurrence allows us to identify the coefficients in closed form and extend the method to every $N \geq 1$ and $q \in (0, 1)$. We also show that this new method satisfies

$$\|x_N - x^\star\| \leq \frac{1+\sqrt{q}}{2\mu}(1-\sqrt{q})^N \|\nabla f(x_0)\| \qquad \text{for all } N \geq 1.$$

References

[1] S. Das Gupta, B. P. G. Van Parys, and E. K. Ryu. Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. Mathematical Programming, 204(1–2):567–639, 2024.

[2] Y. Kamri, J. M. Hendrickx, and F. Glineur. Numerical design of optimized first-order algorithms. arXiv:2507.20773, 2025.

[3] J. Kim, A. Ozdaglar, C. Park, and E. K. Ryu. Time-reversed dissipation induces duality between minimizing gradient norm and function value. In Advances in Neural Information Processing Systems 36 (NeurIPS 2023), 2023.

[4] K. Shu and A. L. Wang. A unified theory of H-duality in first-order methods. arXiv:2609.03281, 2026.

Session 15I: Geometry *Room: Sophie Germain (A1.139)*

Generalized infinite dimensional Alpha-Procrustes based geometries

Salvish Goomanee[*†], Andi Han[‡], Pratik Jawanpuria[§], Bamdev Mishra[¶]

[*]Corresponding author.
[†]CIRB, Collège de France, CNRS, INSERM, Université PSL, Paris, France. Current address: Université Côte d'Azur, INRIA, I3S, Sophia Antipolis, France (salvish.goomanee@inria.fr).
[‡]University of Sydney, Australia and RIKEN AIP, Japan (andi.han@sydney.edu.au).
[§]Indian Institute of Technology Bombay, India (pratik.jawanpuria@iitb.ac.in).
[¶]Microsoft India (bamdevm@microsoft.com).

September 4, 2026

This work extends the recently introduced Alpha-Procrustes family of Riemannian metrics for symmetric positive definite (SPD) matrices by incorporating generalized versions of the Bures-Wasserstein (GBW), Log-Euclidean, and Wasserstein distances. While the AlphaProcrustes framework has unified many classical metrics in both finite and infinite dimensional settings, it previously lacked the structural components necessary to realize these generalized forms. We introduce a formalism based on unitized Hilbert-Schmidt operators and an extended Mahalanobis norm that allows the construction of robust, infinite dimensional generalizations of GBW and Log-Hilbert-Schmidt distances. Our approach also incorporates a learnable regularization parameter that enhances geometric stability in high dimensional comparisons. Preliminary experiments reproducing benchmarks from the literature demonstrate the improved performance of our generalized metrics, particularly in scenarios involving comparisons between datasets of varying dimension and scale. This work lays a theoretical and computational foundation for advancing robust geometric methods in machine learning, statistical inference, and functional data analysis.

Session 15I: Geometry *Room: Sophie Germain (A1.139)*

Topology optimization of PEMFC channels using a deshomogenization method

Luca Gorini[1], Sébastien Kawka[1]

[1]Univ. Grenoble Alpes, CEA, Liten, F-38000 Grenoble, France, luca.gorini@cea.fr

Keywords: topology optimization, homogenization, power source

Proton Exchange Membrane Fuel Cells (PEMFCs) are electrochemical systems that generate electrical energy through the reaction between dihydrogen at anode and dioxygen at cathode. In these systems, dihydrogen (on the anode side) and dioxygen (on the cathode side) flow through microchannels—defined by ribs—and diffuse toward their respective catalytic areas, where the electrochemical reaction occurs. Uniform gas distribution across these surfaces enhances power efficiency, optimizes catalyst utilization and/or limit ageing. In this work, we address the topology optimization of the distribution ribs using a dehomogenization method [1], [2], [3], [4]. The behavior of the gases within the fuel cell is modeled using two isotherm stationnary compressible Stokes systems coupled through a non-linear reaction type boundary term and solved using finite element methods. The ribs are represented as a periodic arrangement of squares, deformed by a diffeomorphism that locally modulates their shape and orientation. We begin by asymptotically deriving a homogenized system, where fine geometric details are embedded in the effective coefficients. The dependence on local dilation and orientation of the deformed periodic medium is explicitly characterized. After an offline computation phase where effective coefficients are precomputed for various values of spacing and orientation parameters of the ribs, an adjoint-based optimisation procedure is applied to determine the optimal geometric arrangements of local rib spacing and orientation. Due to feasibility constraints, ribs at cathode et anode must align correctly. To manage this constraint, established methods for reconstructing ribs and fine channels require adaptation. .

References

[1] Feppon, F., Multiscale Topology Optimization of modulated fluid microchannels based on asymptotic homogenization, Computer Methods in Applied Mechanics and Engineering, 2024.

[2] Dede, E. M. and Zhou, Y. and Nomura, T., Inverse design of microchannel fluid flow networks using Turing pattern dehomogenization, Structural and Multidisciplinary Optimization, 2020.

[3] Wu, J. and Sigmund, O. and Groen, J. P., Topology optimization of multi-scale structures: a review, Structural and Multidisciplinary Optimization, 2021.

[4] Li, H. and Jolivet,P. and Alexandersen, J., Multi-scale topology optimisation of microchannel cooling using a homogenisation-based method, Structural and Multidisciplinary Optimization, 2024.

Session 8G: Optimizing under network structure *Room: Enrico Fermi (A1.133)*

A Branch&Check approach for Selecting Multi-Site Balancing Entities

Emilien Gouffault, Alexandre Marié, Emily Little

RTE R&D, Paris La Défense, France
{emilien.gouffault, alexandre.marie, emily.little}@rte-france.com

Keywords: Mixed-Integer Programming, Optimal Power Flow, Sensitivity Analysis.

In order to preserve frequency stability, Transmission System Operators (TSOs) must maintain generation and load balance at all time. To do this, they rely on reserves, often procured in advance, with varying activation times – some activated automatically and others manually. Since the beginning of 20261 this obligation has been extended to include all renewable energy facilities with an installed capacity above 10 MW. As a result, the french Balancing Mechanism (Méchanisme d'Ajustement) receives a growing volume of offers, which real-time operators increasingly struggle to fully exploit. Moreover, 300+ of these balancing entities are spread across multiple sites, adding a further challenge to processing and analyzing the offers. The core is a branch-and-check decomposition [1]: a master MILP selects and volumeallocates generation offers to meet a signed volume target at minimum cost. Congestion management is enforced through DC load-flow checks that lazily generate linear feasibility cuts based on Power Transfer Distribution Factors (PTDFs), an adaptively-scoped N-1 watchlist, and a final AC validation of the retained operating point. Post-contingency line flows are reconstructed analytically from a single base-case sensitivity computation using Line Outage Distribution Factors (LODF), avoiding a full sensitivity solve per contingency; adaptive electrical neighborhoods, cut-edge exclusion, and low-impact pruning bound the monitored perimeter, and post-incident temporary/permanent (TATL/PATL) loading limits are enforced. Congestion and post-contingency analysis are done using PowSyBl library [2] The formulation is multi-horizon and can handle several types of offers: a single dispatch decision is validated simultaneously across the current situation and forecast horizons to ensure offers' constraints, each with its own network snapshot, sensitivities and contingency set, and supports both upward and downward offers through a unified signed-volume model with typed offer structures (indivisible, prorated, stacked). The tool has been deployed as an interactive web application to be tested by real-time operators. It returns least-cost redispatch, a per-offer / per-area breakdown ranked by marginal price, forecast available margins per horizon, an itemized computation-time decomposition, and voltage-security diagnostics. Results on a national-scale network highlight that security and precision must be traded against computation time. Instead of an exhaustive N-1 analysis over all horizons, the method makes this trade-off explicit through adaptive contingency scoping and perimeter reduction, keeping runtimes compatible with real-time operation.

1Following the implementation of law no2025-391 of 30 April 2025 (DDADUE)(https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000051538879)

References

[1] E. Lam, G. Gange, P. Stuckey, P. Van Hentenryck, and J. Dekker, "Nutmeg: a MIP and CP hybrid solver using branch-and-check," SN Operations Research Forum, 2020.

[2] L. F. Energy. Powsybl (power system blocks), a framework dedicated to the modeling, simulation, and visualization of electrical networks. Available: https://github.com/powsybl

Session 5A: Invited session: Nonconvex landscapes and algorithms in data science *Room: Marie Curie (Auditorium)*

Trust-Region Methods For Burer-Monteiro Factorizations Of Phase Retrieval Problems

Florentin Goyens, Clément W. Royer, Irène Waldspurger

1ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium, goyensflorentin@gmail.com 2LIGM, CNRS, Université Gustave Eiffel, clement.royer@univ-eiffel.fr 3CNRS, DMA ENS, irene.waldspurger@ceremade.dauphine.fr

Keywords: Burer-Monteiro factorization, Phase retrieval, Trust-region algorithms

Factored formulations of semidefinite programs are a popular approach to handle such instances at scale, even though such reformulations worsen the problem conditioning. When the original semidefinite programs arise from phase retrieval problems, additional ill-conditioning may also be present. As a result, first-order algorithms may struggle to converge quickly towards a solution. In this talk, we consider nonconvex factored formulations of phase retrieval semidefinite programs. Using a second-order Riemannian trust-region framework, we establish fast local convergence of the method on a subclass of these problems. In practice, we observe that this fast convergence holds more broadly.

References

[1] F. Fogel, I. Waldspurger, and A. d'Aspremont, Phase retrieval for imaging problems. Math. Program. Comput., 8:311-335, 2016.

[2] F. Goyens and C. W. Royer, Riemannian trust-region methods for strict saddle functions with complexity guarantees. Math. Program., 213:863-905, 2025.

[3] G. Zhang, S. Fattahi, and R. Y. Zhang, Preconditioned gradient descent for overparameterized nonconvex Burer-Monteiro factorization with global optimality certification. J. Mach. Learn. Res., 24:1-55, 2023.

Session 10A: Applied combinatorial optimization *Room: André Marie Ampère (Amphi II)*

Aperiodic Tiling Rhythmic Canons

Gennaro Auricchio[1], Luca Ferrarini[2], Camille Grange[2]

[1]University of Padova, Italy, gennaro.auricchio@math.unipd.it
[2]University Sorbonne Paris Nord, France, {ferrarini, grange}@lipn.univ-paris13.fr

Keywords: Mathematical music, Tiling Rhythmic, Aperiodic Canons, Integer Programming, SAT Encoding

The Aperiodic Tiling Rhythmic Canons Problem (ATCP) is a problem in Mathematical Music Theory that aims at determining all possible aperiodic couples of complementary rhythms for a given period n. More precisely, given n $\in$N, we look for the couples (A, B) $\subseteq$Zn $\times$ Zn of aperiodic sets such that A $\oplus$B = Zn. If we interpret A as the inner voice and B as the outer voice, the tiling condition imposes that the two voices combined give a musical composition such that playing A translated by the elements of B covers all instants time of the bar. Until now, the numerical approaches to this problem were based on having access to at least one of the two elements of the couple, e.g. A, which had to be determined by hand. This problem is intrinsically easier and can be seen as a tiling problem in the cyclic group Zn, with the additional constraint that the complementary voice B is aperiodic. In this ongoing work, we propose a model that shifts the attention back to the set of all feasible couples (A, B), therefore bypassing the construction of A which is usually tedious. We provide to ATCP both an Integer Linear Programming and a SAT formulation. Our contribution thus provides two complementary computational frameworks for the enumeration of all aperiodic tilings rhythmic canons. We present preliminary results on the numerical resolution of this problem with our new models, and validate the benchmarks of the literature for n $\in$\{72, 108, 120, 144, 168\}.

References

[1] Auricchio, G., Ferrarini, L., Lanzarotto, G. (2023). An integer linear programming model for tilings. Journal of Mathematics Music, 17(3), 514–530.

[2] Auricchio, G., Ferrarini, L., Gualandi, S., Lanzarotto, G., Pernazza, L. (2022). A SAT Encoding to Compute Aperiodic Tiling Rhythmic Canons. In: Schaus, P. (eds) Integration of Constraint Programming, Artificial Intelligence, and Operations Research. CPAIOR 2022. Lecture Notes in Computer Science, vol 13292. Springer, Cham.

[3] Auricchio G., Ferrarini L., Gualandi S., Lanzarotto G., Pernazza L., Computing Aperiodic Tiling Rhythmic Canons via SAT Models. Constraints 29, 215–233 (2024).

[4] Derval, G., Lecoutre, C. (2026). Aperiodic Tiling and Rhythmic Canons: A CP Journey. Proceedings of the AAAI Conference on Artificial Intelligence, 40(17), 14209–14216.

Session 12K: Algorithms 2.0 : advances in computational methods *Room: Loggia (A1 - first floor)*

EvoHGS: Co-evolving Language Models with Hybrid Genetic Search for Vehicle Routing Problems

Zihao Guo, Yiding Shi, Xia Jiang, Yaoxin Wu, Jie Gao, Andrea Araldo, Mounîm El-Yacoubi

1Institut Polytechnique de Paris, France, ip-paris.fr 2Nanyang Technological University, Singapore, ntu.edu.sg 3Eindhoven University of Technology, The Netherlands, tue.nl 4Delft University of Technology, The Netherlands, tudelft.nl 5Télécom SudParis, France, telecom-sudparis.fr

Keywords: Vehicle routing problems, hybrid genetic search, large language models, evolutionary computation, reinforcement learning

Vehicle Routing Problems (VRPs) are a fundamental class of combinatorial optimization problems aimed to find high-quality routes for vehicles under operational constraints. Despite decades of study in operations research, VRPs remain computationally challenging due to their NP-hardness, with additional difficulty introduced by the diverse constraint structures that arise across different variants. Hybrid Genetic Search (HGS) has emerged as a state-of-the-art metaheuristic framework for many VRP variants. However, its performance still depends heavily on manually designed search operators and empirical parameter tuning, limiting its ability to adapt automatically to complex and diverse problem settings. To address the limitation, we propose EvoHGS, a large language model (LLM)-guided evolutionary framework that automatically improves HGS across the VRP family. EvoHGS integrates LLM-based heuristic generation with evolutionary search, allowing candidate HGS operators to be generated, evaluated, selected, and refined through an iterative optimization process. To further align generation with downstream search performance, EvoHGS uses feedback collected during evolution to finetune the LLM via reinforcement learning. This enables the model to progressively produce operators that are better adapted to the search dynamics of HGS across VRP variants. Experiments on multiple VRP variants show that EvoHGS consistently outperforms the original HGS baseline at the scale used for evolution, remains competitive when the evolved operators are applied to larger instances, and demonstrates promising zero-shot generalization to unseen VRP variants.

Session 9A: Physics Informed Neural Networks *Room: Marie Curie (Auditorium)*

Parametric Physics Informed Neural Network for the Finite Horizon classical Lady-in-the-Lake Pursuit-Evasion Differential Game

Dhruv Gupta

L2S, Centralesupelec, Universite Paris Sacaly

Keywords: Pursuit-evasion differential games, Lady-in-the-Lake problem, Hamilton-Jacobi-Isaacs equations, Physics-Informed Neural Networks, Parametric PINNs

We consider a finite-horizon classical Lady-in-the-Lake pursuit-evasion differential game in which an evading agent moves inside a circular region with velocity boundμ, while a boundary-constrained pursuer seeks interception. The game is formulated through a Hamilton-Jacobi-Isaacs (HJI) equation, whose repeated numerical solution can become computationally demanding when the game parameters change. We develop a Physics-Informed Neural Network (PINN) formulation that approximates the HJI value function by combining the viscous HJI residual with dynamic-programming consistency. The proposed framework is then extended to a parametric PINN in which the Lady's velocity$\mu$and terminal horizon$T$are provided as network inputs, allowing a single trained model to represent a continuous family of game configurations. The learned solutions are validated against a classical grid based Semi-Lagrangian numerical solver for fixed and previously unsampled parameter combinations. Finally, we consider an adaptive execution scenario in which the Lady's velocity changes unexpectedly during the game. Rather than retraining the network or performing a new numerical HJI solve, the updated feedback strategy is obtained by direct evaluation of the trained parametric PINN using the new velocity and remaining horizon. Numerical experiments for both one-pursuer and two-pursuer configurations demonstrate the accuracy of the learned value functions and illustrate the potential of parametric PINNs for rapid policy adjustment in changing pursuit-evasion environments.

Session 12J: Handling dynamics *Room: Emmy Noether (A1.140)*

Superposition Principles for Perturbed Sweeping Processes and Differential Inclusions

Aldo Gutiérrez-Concha, Francisco Silva, Emilio Vilches

1Departamento de Ingeniería Matemática, Universidad de Chile, Chile, agutierrez@dim.uchile.cl 2 Laboratoire XLIM, Université de Limoges, France, francisco.silva@unilim.fr 3 Instituto de Ciencias de la Ingeniería, Universidad de O'Higgins, Chile, emilio.vilches@uoh.cl 4 Centro de Modelamiento Matemático, Universidad de Chile, Chile

Keywords: sweeping process, superposition principle, optimal transport, differential inclusion, optimal control

We establish a superposition principle for perturbed sweeping processes governed by a moving prox-regular set. The classical Superposition Principle connects the Eulerian description of dynamics through the continuity equation with their Lagrangian description in terms of characteristic trajectories. Our result extends this framework to constrained dynamics and, in the unconstrained case, extends the superposition principle for differential inclusions of Cavagnari, Marigonda, and Piccoli [1] to non-autonomous dynamics that may be discontinuous in time. More precisely, starting from a probability measure on path space concentrated on Eulerian admissible sweeping trajectories, we show that its time marginals form a narrowly Lipschitz continuous curve of probability measures supported on the moving constraint set and solving a constrained continuity equation driven by an admissible Borel velocity field. Conversely, every solution of this continuity equation admits a superposition representation concentrated on Eulerian admissible sweeping trajectories. This yields a rigorous correspondence between the Lagrangian and Eulerian descriptions of the dynamics. We also prove a discrete version of the principle, showing that admissible sweeping trajectories can be approximated in Wasserstein distance by empirical measures supported on finitely many characteristics. As an application, we consider an optimal control problem governed by the perturbed sweeping process and prove, under a convexity assumption on the running cost, that its Kantorovich and Eulerian formulations have the same value and both admit minimizers. This equivalence is in the spirit of the multi-agent control framework of Cavagnari, Lisini, Orrieri, and Savaré [2].

References

[1] G. Cavagnari, A. Marigonda, B. Piccoli, Superposition principle for differential inclusions, Lecture Notes in Computer Science 10665 (2018), 201–209.

[2] G. Cavagnari, S. Lisini, C. Orrieri, G. Savaré, Lagrangian, Eulerian and Kantorovich formulations of multi-agent optimal control problems: equivalence and Γ-convergence, J. Differential Equations 322 (2022), 268–364, arXiv:2011.07117.

Session 15F: Combinatorial Optimization & Applications *Room: Ada Lovelace (A1.128)*

Drawing and Scheduling Incomplete Tournaments with Seeding Constraints: the Case of the UEFA Champions League

Julien Guyon, Adle Ben Salem, Thomas Buchholtzer, Mathieu Tanré

1CERMICS, CNRS, ENPC, IP Paris, France, {julien.guyon, adle.ben-salem, thomas.buchholtzer, mathieu.tanré}@enpc.fr
2MathRisk, Inria, Paris, France 3Department of Finance and Risk Engineering, NYU Tandon, USA

Keywords: OR in sports, fairness in sports, sports scheduling, graph theory, integer programming

We investigate and compare four draw procedures for incomplete round-robin tournaments with p seeding pots and q teams per pot, where each team faces exactly k different opponents from each pot. First, we consider a method where matchups are drawn before the match schedule is built. We show, using a graph-theoretical argument, that, if $q = s(k + 1)$ with $s \geq 3$, the scheduling issue cannot be ignored when drawing the matchups, by exhibiting a draw outcome that cannot be scheduled within pk match days. Second, we study an alternative method where one first builds a schedule template satisfying all the pot constraints before randomly populating it with the pq teams. We study the overall minimum number of breaks achievable by the schedule template, and in the case of the UEFA Champions League (UCL), we build a template that minimizes the number of breaks and optimizes various schedule fairness and attractiveness criteria. For both methods, a randomized variant is also considered. Using Monte Carlo simulations, we assess opponent strength fairness for the four methods in the case of the 2024–25 UCL by comparing the distributions of average opponent strength that they produce for all the teams. We find that all four procedures are fair to all teams—no team is significantly aggrieved—and the first procedure yields more uniform matchup probabilities than the second one. Finally, we introduce a luck index that objectively ranks teams from the luckiest to the unluckiest during the actual draw.

References

[1] J. Guyon, A. Ben Salem, T. Buchholtzer, M. Tanré, Drawing and scheduling incomplete tournaments with seeding constraints: the case of the UEFA Champions League, doi:10.2139/ssrn.5413142.

Session 12G: Mid and Long Term Energy Management *Room: Enrico Fermi (A1.133)*

Coupled Investments in SMRs and Industrial Sites: a Real Options Analysis

Cyrian Hallermeyer[1,2,3]

[1]SERG, LGI CentraleSupélec, France
[2]EDF R&D, France
[2]IFP School, France
cyrian.hallermeyer@centralesupelec.fr

Keywords: Real options, Risk management, Game Theory, Industrial heat, Decarbonisation

Decarbonising industrial heat has become a pivotal challenge for policymakers and industrial stakeholders alike. However, investment in low-carbon heat technologies is still timid. This is hindered, among other factors, by risk aversion from economic agents having to take investment decisions, the irreversible nature of these investments and the high sunk costs they would involve if future outcomes were unfavourable. As a result, firms may delay or avoid investments, waiting for greater certainty in energy prices, regulatory frameworks, or market signals [1]. In the context of the development of innovative, decarbonised heat production technologies, Small Modular Reactors (SMRs) are gaining attention as potential solutions for providing low-carbon, mediumto-high temperature heat to industrial sites [2]. This paper explores the potential role of Small Modular Reactors (SMRs) in supplying decarbonised heat and electricity to industrial sites. In particular, we explore the possible emergence of bilateral contracts between SMR developers and industrial firms, which may offer a viable pathway to overcome investment hurdles and align supply with demand. To investigate this, we develop a coupled investment model within a Real Options Analysis framework, capturing the strategic decision-making processes for both an industrial firm and a SMR developer. We combine Dynamic Programming with Game Theory to solve for the strategic behavior of both firms. The model considers multiple investment choices for the industrial firm, incorporating the irreversibility of capital expenditures and the potential for bilateral contracting in heat and electricity supply with the SMR firm. We solve the model by backward induction, adapting the Least-Square Monte Carlo method [3] to a two-agent framework in a risk averse fashion. Risk aversion is modelled using the Expected conditional risk measures [4]. Our results suggest that both firms delay their investment decisions when given the option to wait before investing: The industrial firm keeps its initial gas boiler for a longer period, while the SMR firm holds off investment and construction of the SMR. Furthermore, risk aversion exacerbates the postponing of investments in non-nuclear investments. Finally, when contracting between firms is allowed, an increased risk aversion motivates both firms' to sign heat supply contracts in order to hedge against energy price risks.

References

[1] A. K. Dixit and R. S. Pindyck. Investment Under Uncertainty. Princeton University Press, 1994.

[2] NEA and OECD. Beyond Electricity: The Economics of Nuclear Cogeneration. OECD, 2022.

[3] F. A. Longstaff and E. S. Schwartz. Valuing American Options by Simulation: A Simple Least-Squares Approach. Review of Financial Studies 14.1 (2001), 113–147.

[4] T. Homem-de-Mello and B. K. Pagnoncelli. Risk aversion in multistage stochastic programming: A modeling and algorithmic perspective. European Journal of Operational Research 249.1 (2016), 188–199.

Session 8E: Optimisation with chance constraints *Room: Antoine Lavoisier (A1.122)*

Chance-constrained Linear Complementarity Problems

René Henrion, Martin Schmidt

[1]Weierstrass Institut Berlin, Germany; [2]Trier University, Germany

Keywords: Linear Complementarity Problems, Chance Constraints, Existence of Solutions, Convexity, Differentiability, Optimality Conditions

We study linear complementarity problems (LCPs) under uncertainty, which we model using chance constraints. Since the complementarity condition of the LCP is an equality constraint, it is required to consider relaxations, which naturally leads to optimization problems in which the relaxation parameters are minimized for given probability levels. We focus on these optimization problems and first study the continuity of the related probability functions and the compact- ness of the feasible sets. This leads to existence results for both types of models: one with a joint chance constraint and one with separate chance constraints for both uncertainty-affected conditions of the LCP. For both, we prove the differentiability of all probability functions and derive respective gradient formulae. For the separate case, we prove convexity of the respective optimization problem and use the gradient formulae to derive necessary and su cient optimal- ity conditions. In a small example of a Cournot oligopoly among energy producers, we nally illustrate our theoretical findings.

Session 12B: Invited session: Discrete optimization for hydro unit commitment problem *Room: Henri Poincaré (Amphi I)*

Piecewise linear approximation for the hydro unit commitment problem

Kea Horvath, Christian Artigues, Pascale Bendotti, Sandra Ulrich Ngueveu, Cécile Rottner

[1]LAAS-CNRS, University of Toulouse, CNRS, INP, Toulouse, France; [2]EDF Lab Paris-Saclay, Palaiseau, France

Keywords: single-plant hydro unit commitment, piecewise linear approximation, mixed-integer linear programming

The single-plant hydro unit commtment problem (1-HUC) involves determining the turbine flow rate at each time step over a discretized horizon in order to maximize production value. One of the challenges lies in the power function –which is neither linear nor concave– requiring approximation. Currently, EDF discretizes this function into a finite set of operating points (PDF), significantly limiting the range of obtainable solutions. This internship investigates the alternative of piecewise linear (PWL) approximation, which allows for continuous flow rates within each segment. In [1], it was established that PWL approximation yielded promising results compared to PDF and non-linear formulations but a single PWL formulation was used in a limited experimental setting. We proposed five mixed-integer linear programming models for this approximation. They encode the same linearization and differ only in the method used to select the active segment: one binary variable per segment, a logarithmic encoding of these variables, one variable per concave piece, an incremental formulation requiring integrality only for the segments that initiate a piece, and Gurobi's native piecewise linear constraint. The best-performing model was then compared against the PDF model using an equivalent approximation budget – that is, matching the number of operating points to the number of segments in the linearization. As existing problem instances could not be directly applied to our models, we also developed a generator for realistic and varied instances. Experiments conducted on 5,200 configurations show that the incremental formulation out- performs the others. For an equivalent budget, the PWL approach yields better solutions than the PDF approach while also being faster. Piecewise linear approximation therefore represents a viable alternative to operating-point discretization for the 1-HUC problem.

References

[1] A. Heintzmann, C. Artigues, P. Bendotti, S. U. Ngueveu and C. Rottner. A comparison of alternative models for solving a non-linear single plant Hydro Unit Commitment prob- lem Computers & Operations Research, 165, 2024 https://doi.org/10.1016/j.cor.2024. 106591.

Session 13B: On complexity *Room: Ada Lovelace (A1.128)*

On the Complexity of some Unicity Problems

Olivier Hudry

Télécom Paris, France, olivier.hudry@telecom-paris.fr

Keywords: Complexity Theory, Combinatorial Optimisation, Unicity of Solution, UScompleteness, Satisfiability Problem

In the theory of complexity, decision problems are problems stating a question that admits only the answer "Yes" or "No". The question can be set in the very general following form: "Given an instance I and a property P on I, is P true for I?", where P can be expressed as: "Is there an appropriate entity (which can be a set, a function, a number...; such an entity depends on the problem) satisfying a given characteristic?". For instance, for the well-known problem Satisfiability, the instance I is given by a Boolean formula F defined on a finite set S of Boolean variables and the considered property is the question: "Is there an assignment function defined on S allowing to satisfy F, i.e. providing the Boolean value True to F?". In this communication, we consider the problems obtained by adding the extra word "unique" to the usual question: "Is there a unique entity satisfying the given characteristic?". Thus, for the problem Satisfiability, we obtain the problem called Unique Satisfiability in [1], stated as follows: "Given a Boolean formula F defined on S, is there a unique assignment function defined on S allowing to satisfy F?". We study the complexity of several usual problems when we consider the question of the unicity of the searched entity. More precisely, we intend to compare the complexity of the existence of a solution to the complexity of the unicity of such a solution. Some problems are polynomial for both versions. It is the case for instance for the decision problem associated with the shortest path problem in positively-weighted graphs, or for the decision problem associated with the minimum spanning tree problem. For some other problems, under appropriate hypotheses relative to the complexity classes (like the usual hypothesis saying that P and NP are different), the two versions do not belong to the same class. For instance, Satisfiability belongs to NP (as known, Satisfiability is even NP-complete) while Unique Satisfiability is US-complete (where "US" stands for the "Unique Solution" class) and thus does not belong to NP (at least if NP and co-NP are distinct, since US contains co-NP). Similarly, the other famous problem Hamiltonian Cycle belongs to NP (it is also an NP-complete problem) while the version in which we consider the unicity of the Hamiltonian cycle is US-complete ([2]). The aim of the talk is to compare the complexities of such problems when taking the unicity property into account.

References

[1] A. Blass, Y. Gurevich, On the unique satisfiability problem, Information and Control, 55: 80-88, 1982.

[2] O. Hudry, A. Lobstein, On the complexity of determining whether there is a unique Hamiltonian cycle in a graph, WSEAS Transactions on Mathematics, 21: 433–446, 2022.

Session 15B: Learning & Neural Networks *Room: Henri Poincaré (Amphi I)*

Landscape Analysis for Shallow Neural Networks: Complete Classification of Critical Points for Cubic Activation and Affine Target Functions

Shokhrukh Ibragimov[1], Ilkhom Mukhammadiev[2], Diyora Salimova[3]

[1]Applied Mathematics: Institute for Analysis and Numerics, University of Münster, Germany; sibragim@uni-muenster.de
[2]Department for Applied Mathematics, University of Freiburg, Germany; ilkhom.mukhammadiev@mathematik.uni-freiburg.de
[3]Department for Applied Mathematics, University of Freiburg, Germany; diyora.salimova@mathematik.uni-freiburg.de

Keywords: shallow neural networks, polynomial neural networks, optimization landscape, nonglobal local minima, critical points

We study the optimization landscape induced by the true loss for shallow polynomial neural networks (PNNs) with h ∈N neurons in the hidden layer, one-dimensional input and output layers, and a monomial activation function of degree d ∈N, trained against a non-constant affine linear target function. Our first main result provides, for arbitrary activation degree d, a sharp existence and nonexistence criterion for global minimizers together with necessary structural conditions. We show that the infimum of the loss is always zero and is attainable with at least d active and visible hidden neurons, that is, hidden neurons with non-zero inner and outer weights, having pairwise distinct pivots. In contrast, if h < d, then the infimum cannot be attained, and every minimizing sequence of parameters necessarily diverges to infinity. Our second main result provides a complete classification of all critical points of the loss function for cubic activation. We show that the loss landscape admits no local maximizers, that critical points cannot have exactly two distinct pivots, and that global minimizers require at least three distinct pivots. Critical points with no active hidden neurons correspond only to saddle points. Consequently, non-global local minimizers and non-trivial saddle points arise only in networks where all pivots coincide. Moreover, non-global local minimizers require all hidden neurons to be active and visible, with exactly one hidden neuron having a slope sign matching that of the target function. We further show that each hidden neuron of a critical point that is not a global minimizer has either an inputdependent or zero contribution, but no nonzero input-independent contribution, to its corresponding realization function.

Session 8K: Invited session : Quantum Applications in Operations Research part 2 *Room: Loggia (A1 - first floor)*

Quantum Genetic Algorithm for Combinatorial Optimisation

Adil Kassaoui, Eric Bourreau

Abstract

As one of the largest postal operators, La Poste is directly confronted with large-scale combinatorial optimization problems, particularly variants of the Vehicle Routing Problem (VRP). Among the best- performing heuristics for such problems is Hybrid Genetic Search (HGS), which — like many state-of-the- art combinatorial optimization methods — combines a genetic algorithm with local search procedures, following the general paradigm known as memetic algorithms. Beyond these classical approaches, quantum computing has emerged as a new perspective on combinatorial optimization. Current hardware, however, remains in the Noisy Intermediate-Scale Quantum (NISQ) era, where only a small number of physical qubits are available. This constraint motivates the use of LogQ encoding, which represents solutions using exponentially fewer qubits than direct encodings, making it compatible with near-future hardware — but at the cost of significantly restricting the freedom available in designing quantum encodings and algorithmic operations. Genetic algorithms are naturally suited to this constraint, as they operate through generic variation and selection operators rather than requiring rich circuit structures or gradient access. Combined with the strong classical track record of genetic approaches such as HGS, this makes the genetic paradigm a well-motivated candidate for building an optimization algorithm within the LogQ framework.

Résumé

En tant que l'un des plus grands opérateurs postaux, La Poste est directement confrontée à des problèmes d'optimisation combinatoire à grande échelle, en particulier des variantes du problème de tournées de véhicules (Vehicle Routing Problem, VRP). Parmi les heuristiques les plus performantes pour ce type de problèmes figure la recherche génétique hybride (Hybrid Genetic Search, HGS), qui — comme de nombreuses méthodes d'optimisation combinatoire de référence — combine un algorithme génétique à des procédures de recherche locale, selon le paradigme général dit mémétique. Au-delà de ces approches purement classiques, l'informatique quantique a émergé comme une nouvelle perspective pour l'optimisation combinatoire. Le matériel actuel reste néanmoins dans l'ère NISQ (Noisy Intermediate- Scale Quantum), où seul un nombre restreint de qubits physiques est disponible. Cette contrainte motive l'usage de l'encodage LogQ, qui représente les solutions à l'aide d'un nombre de qubits exponentiellement plus faible que les encodages directs, le rendant compatible avec du matériel à court terme — mais au prix d'une réduction significative de la liberté offerte dans la conception des encodages quantiques et des opérations algorithmiques. Les algorithmes génétiques s'adaptent naturellement à cette contrainte, puisqu'ils reposent sur des opérateurs génériques de variation et de sélection plutôt que sur des structures de circuits riches ou un accès au gradient. Associé aux bonnes performances classiques déjà établies des approches génétiques telles que HGS, cela fait du paradigme génétique un candidat bien motivé pour la construction d'un algorithme d'optimisation dans le cadre de l'encodage LogQ.

Session 12E: Convergence rates *Room: Antoine Lavoisier (A1.122)*

Convergence Rates for Stochastic Proximal Methods via Regularized Optimality Gaps

Lucas Ketels, Daniel Cortild, Juan Peypouquet, Guillaume Garrigos

1Université Paris Cité and Sorbonne Université, CNRS, {ketels,garrigos}@lpsm.paris 2University of Oxford, Mathematical Institute, daniel.cortild@maths.ox.ac.uk 3University of Groningen, Bernoulli Institute, j.peypouquet@rug.nl,

Keywords: Stochastic proximal point, stochastic forward-backward, convex optimization

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 [3, 4], or smoothness, less commonly used in the SPP literature. Existing results [1] 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 (in the fashion of [1, 2]). 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.

References

[1] Javier I. Madariaga, Convergence of the Iterates of the Stochastic Proximal Gradient Method, arXiv:2604.13388,

[2] Patrick L. Combettes, Jean-Christophe Pesquet, Stochastic Approximations and Perturbations in Forward-Backward Splitting for Monotone Operators, arXiv:1507.07095,

[3] Hilal Asi, John C. Duchi, Stochastic (Approximate) Proximal Point Methods: Convergence, Optimality, and Adaptivity, SIAM Journal on Optimization, arXiv:1810.05633

[4] Andrei Patrascu, Ion Necoara, Nonasymptotic convergence of stochastic proximal point methods for constrained convex optimization, arXiv:1706.06297.

Session information not available in the current online program.

Variational Quantum Algorithms for the Uncapacitated Single-Allocation p-Hub Median Problem

Rashed Khanjani-Shiraz[1]

[1]Technical University of Munich, Germany, Rashed.Khanjani-Shiraz@tum.de

Keywords: hub location, QUBO, variational quantum algorithms

This work develops a hybrid quantum-classical framework based on variational quantum algorithms for the uncapacitated single-allocation p-hub median problem, a classical hub-location problem in which hub-selection and node-allocation decisions are coupled [1]. We first analyze its binary structure and derive a quadratic representation of the transportation objective. The hub-cardinality, single-allocation, and hub-activation constraints are incorporated through quadratic penalty terms, leading to a Quadratic Unconstrained Binary Optimization (QUBO) representation suitable for an Ising-Hamiltonian encoding [2]. Building on this formulation, we focus on hybrid variational quantum algorithms. In the resulting framework, a parameterized quantum circuit prepares candidate binary configurations, measurements are used to estimate the corresponding QUBO objective, and a classical optimizer iteratively updates the circuit parameters. We consider the Variational Quantum Eigensolver (VQE) [3] and the Quantum Approximate Optimization Algorithm (QAOA) [4] as two representative variational approaches for this formulation. Particular attention is given to the encoding of the hub-selection and allocation variables, the construction of the penalized objective, and the role of penalty parameters in maintaining feasibility. The aim is to establish a transparent QUBO-based framework for applying variational quantum algorithms to the p-hub median problem, providing a basis for subsequent computational experiments comparing circuit architectures, classical parameter optimizers, penalty choices, and VQE/QAOA behavior on benchmark hub-location instances.

References

[1] M. E. O'Kelly, A quadratic integer program for the location of interacting hub facilities, European Journal of Operational Research 32(3) (1987), 393–404, doi:10.1016/S0377-2217(87)80007-3.

[2] A. Lucas, Ising formulations of many NP problems, Frontiers in Physics 2 (2014), 5, doi:10.3389/fphy.2014.00005.

[3] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nature Communications 5 (2014), 4213, doi:10.1038/ncomms5213.

[4] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm, arXiv:1411.4028, 2014.

Session 12A: Invited session: Applications of optimization to power systems and networks *Room: Marie Curie (Auditorium)*

Choice-Based Dynamic Pricing with Deferred Departures in Bike-Sharing Networks

Mesut Can Köseoğlu, Shadi Sharif Azadeh, Luce Brotcorne

1Centre INRIA de l'Université de Lille, Team INOCS, Lille, France, mesut-can.koseoglu@inria.fr, luce.brotcorne@inria.fr
2Department of Transport and Planning, Delft University of Technology, Delft, The Netherlands, S.SharifAzadeh@tudelft.nl

Keywords: Stochastic Programming, Pricing, Bike Sharing, Rolling Horizon

Bike-sharing systems (BSS) are among the most widely adopted new shared mobility services. Because they rely on one-way trips, changes in demand cause imbalances, with bikes depleted in some locations and overflowing in others. Traditionally, this challenge is addressed through costly truck-based rebalancing operations. In high-traffic zones or during peak time periods, precisely when reliability is most critical, operators face a fundamental challenge between service and inventory management, which in station-based BSS, increases further with capacity restrictions. To improve the reliability of the service, BSS operators have introduced reservation options that guarantee bike availability at specific stations and times, which comes at a trade-off in service levels. Reserving bikes too aggressively reduces system utilization while exposing the operator to the risk of not fulfilling future commitments. In this study, we address a choice-based dynamic pricing problem for demand management in BSS networks with a service-level-maximizing objective. We propose a novel multistage stochastic program (MSSP) in which the operator offers a menu of alternatives for each origin station, including destination stations, deferred departures, and prices. The users' choices are integrated through a multinomial choice model, where they either choose a trip alternative from the offered menu or opt out. We also include exogenous pre-booked trips; hence, the operator must decide based on residual capacities. We propose a rolling horizon algorithm to reduce the size of the MSSP by combining a stochastic decision window and expected value approximation of the remaining horizon in a single optimization model [1, 2], enabling menu decisions that adapt to the observed network state. Furthermore, we develop a station-wise Lagrangian decomposition for scalability. The proposed model and solution algorithms are evaluated on a synthetic network with instances diversified in size and demand patterns.

Kaynaklar

[1] G. Pantuso and T. K. Boomsma, "On the number of stages in multistage stochastic programs," Annals of Operations Research, vol. 292, no. 2, pp. 581–603, 2020.

[2] F. Maggioni and G. C. Pflug, "Bounds and approximations for multistage stochastic programs," SIAM Journal on Optimization, vol. 26, no. 1, pp. 831–855, 2016.

Session 15D: Optimisation challenges in modern mobility *Room: Lise Meitner (A1.116)*

Feasibility, Cost, and Environmental Assessment of Heavy Duty Truck Fleet Electrification via Mixed Integer Optimization

Malak Al Kostit[1,2], Sophie Demassey[1], Valentina Sessa[1], Franck Tirard[2], Pierre Chaufour[2]

[1]Centre for Applied Mathematics (CMA), Mines Paris PSL, France, malak.al_kostit@minesparis.psl.eu
[2]Renault Trucks, Lyon, France

Keywords: Mixed Integer Linear Programming, Battery electric trucks, Total cost of ownership, Life cycle assessment, Greenhouse gas emissions.

Road freight transport is a critical target of European decarbonization policy: heavy-duty trucks in particular are subject to a regulatory trajectory mandating a 90% reduction by 2040 relative to 2019 levels of CO2 emissions. Battery electric trucks (BETs) are a leading decarbonization pathway for the sector [1], yet fleet operators lack decision tools that jointly account for the physical feasibility of electrification, given real vehicle usage, and its economic and environmental consequences. This work develops a data-driven optimization framework, operated by a European manufacturer, for the electrification of heavy-duty truck fleets, applied to real client fleets. Raw telematics pings are processed through a pipeline that reconstructs missions, along which potential charging opportunities are identified, from which battery state of charge trajectories are simulated for a catalog of candidate BETs' models. Building on these simulations, we formulate mixed integer linear programs assigning BET models to vehicles. We then perform an a posteriori assessment, comparing the initial all-diesel fleet against each optimized electrified configuration along three criteria: total cost of ownership (TCO) [2, 3, 4], life cycle greenhouse gas emissions [5] and criticality of materials use through an antimony equivalent indicator. This framework grounds fleet electrification decisions in real operational data, with cost and environmental criteria to be integrated directly into the optimization objective in future work.

References

[1] W. Shoman et al., Battery electric long-haul trucks in Europe: Public charging, energy, and power requirements, Transp. Res. D 121, 103825, 2023.

[2] B. Noll et al., Analyzing the competitiveness of low-carbon drive-technologies in roadfreight: A total cost of ownership analysis in Europe, Appl. Energy 306, 118079, 2022.

[3] Z. Wang et al., A total cost of ownership analysis of zero emission powertrain solutions for the heavy goods vehicle sector, Journal of Cleaner Production 434, 139910, 2024.

[4] T. Costantino et al., Standardized framework for techno-economic assessment of heavy-duty vehicles: Insights from a state-of-the-art review, Appl. Energy 415, 127883, 2026.

[5] T. Zhou et al., Life cycle GHG emissions and lifetime costs of medium-duty diesel and battery electric trucks in Toronto, Canada, Transp. Res. D 55, 91–98, 2017.

Session 12E: Convergence rates *Room: Antoine Lavoisier (A1.122)*

Beyond Softmax and Entropy: Convergence Rates of Policy Gradients with f-SoftArgmax Parameterization & Coupled Regularization

Safwan Labbi[1], Daniil Tiapkin[1,2], Paul Mangold[1], Eric Moulines[3,4]

Published as a conference paper at ICLR 2026

[1]CMAP, CNRS, École Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France
[2]Université Paris-Saclay, CNRS, LMO, 91405 Orsay, France
[3]Mohamed bin Zayed University of Artificial Intelligence, UAE
[4]LRE EPITA, 94270 Le Kremlin-Bicêtre, France
{safwan.labbi, daniil.tiapkin, paul.mangold}@polytechnique.edu; eric.moulines@mbzuai.ac.ae

Policy gradient methods are known to be highly sensitive to the choice of policy parameterization. In particular, the widely used softmax parameterization can induce ill-conditioned optimization landscapes and lead to exponentially slow convergence. Although this can be mitigated by preconditioning, this solution is often computationally expensive. Instead, we propose replacing the softmax with an alternative family of policy parameterizations based on the generalized f-softargmax. We further advocate coupling this parameterization with a regularizer induced by the same f-divergence, which improves the optimization landscape and ensures that the resulting regularized objective satisfies a Polyak– Łojasiewicz inequality. Leveraging this structure, we establish the first explicit non-asymptotic last-iterate convergence guarantees for stochastic policy gradient methods for finite MDPs without any form of preconditioning. We also derive sample-complexity bounds for the unregularized problem and show that f-PG with Tsallis divergences achieves polynomial sample complexity in contrast to the exponential complexity incurred by the standard softmax parameterization.

Session 12B: Invited session: Discrete optimization for hydro unit commitment problem *Room: Henri Poincaré (Amphi I)*

Target-Volume Formulations in the Hydro Unit Commitment Problem

Marguerite Lafeuillade, Sourour Elloumi, Safia Kedad-Sidhoum, Cécile Rottner

1 CEDRIC - CNAM 2 UMA - ENSTA 3 EDF

Keywords: Hydro Unit Commitment, Combinatorial optimization, Decomposition approach

The Hydro Unit Commitment (HUC) problem is a difficult sub-component of the Unit Commitment problem (UCP) in electricity production planning. We consider a reservoir valley in which several hydroelectric power plants are interconnected. The objective of the HUC is to optimize the planning of the hydroelectric units in order to satisfy the demand while maximizing the profit and respecting the operational constraints. Its combinatorial complexity as well as its strict volume constraints make it particularly difficult and real large-scale instances remain hard to solve in a satisfactory way. These problems have so far mostly been modeled using linear optimization with integer variables, which remains the dominant approach in the literature [Taktak and D'Ambrosio, 2017]. We focus here on the combinatorial structure induced by the discrete operating points and the target volume constraints of the reservoirs. When the initial and final volumes of a reservoir are fixed, they determine a target volume to be turbined by each plant over the horizon. Since plants operate on a discrete set of cumulative flows, reaching this target constrains the possible combinations of operating points over time. We propose a new extended formulation based on these target volumes, using cardinality variables to describe how many periods are spent at each operating point. We first investigate a preprocessing approach that enumerates all combinations compatible with the given target. We then propose a decomposition scheme that avoids full enumeration by alternating between two mixed-integer programs. One relaxes operating variables to continuous values while keeping these cardinality variables integer, generating candidate combinations with an upper bound on their value. The other fixes a candidate combination and checks whether it can be realized by an integer-feasible schedule. The combination is then excluded from subsequent iterations via a no-good cut, in the spirit of combinatorial Benders decompositions [Côté et al., 2014], until optimality is proven. These approaches proved to be very effective and led to significant improvements in computation times on our set of instances.

Acknowledgment: This research benefited from the support of the FMJH Program Gaspard Monge for optimization and operations research and their interactions with data science.

References

[Côté et al., 2014] Côté, J.-F., Dell'Amico, M., and Iori, M. (2014). Combinatorial benders' cuts for the strip packing problem. Operations Research, pages 643–661.

[Taktak and D'Ambrosio, 2017] Taktak, R. and D'Ambrosio, C. (2017). An overview on mathematical programming approaches for the deterministic unit commitment problem in hydro valleys. Energy systems, pages 57–79.

Session 15C: Energy Systems *Room: André Marie Ampère (Amphi II)*

Optimization of the energy consumption of buildings

Michael Laporte[1], Claudia D'Ambrosio[2], Frederic Meunier[1]

[1]Cermics, France, michael.laporte@enpc.fr, frederic.meunier@enpc.fr
[2]LIX, France, dambrosio@lix.polytechnique.fr

Keywords: Building energy management, geothermal storage, stochastic optimal control, simulation-based policy evaluation, network optimization

Buildings account for 15–20% of global CO2 emissions, with heating and cooling alone responsible for roughly 8%. This work studies how combining photovoltaic (PV) production with geothermal storage can reduce this footprint: PV output is used first to cover a building's immediate electricity needs, and any surplus is stored as heat several tens of meters underground, through a network of buried pipes, to be recovered months later – typically in winter – to preheat the building and cut electricity demand for heating. This principle is being tested at full scale on four residential buildings equipped by Energy4Climate (E4C) on the Saclay plateau, where replacing a large share of the gas supply with an optimized electrical management strategy could cut CO2 emissions by up to 78%. Operating such a system well is nonetheless a genuinely hard control problem: at every instant, the building must arbitrate between drawing on PV, discharging thermal storage, and buying grid electricity, under uncertainty on sunshine, occupant demand, and prices – a decision that must also be reconciled with the far slower, months-long dynamics of the ground storage itself, whose state today constrains what will be available the following winter. The demonstrator will later be extended with electric-vehicle charging stations, adding both further flexibility and further uncertainty, as vehicles double as consumption points and as mobile, decentralized storage whose availability depends on user behavior. A first contribution of this work is a time-stepped simulator of such hybrid PV, geothermal, and gas building energy systems. Both the network topology (which components are present and how they are connected) and the operating scenario (weather, prices, carbon intensity) are fully user-defined and independent of one another, so the same tool can evaluate a candidate control policy – a simple heuristic or an externally supplied, time-varying strategy – across many different building configurations and many realistic, possibly multi-year, conditions imported directly from public meteorological and grid-operator data. It reports the resulting energy consumption, cost, carbon footprint, and comfort, already providing reliable and realistic metrics to compare candidate strategies. The second contribution is ongoing. We are currently developing simplified and tractable versions of the model of the operation problem described above, together with efficient solution methods. These methods will then be evaluated in the simulator to assess their expected performance in practice. This model will serve as a basis to determine how much precision the predictions and the decisions actually require.

Session 8I: Decision making under uncertainty *Room: Sophie Germain (A1.139)*

What Makes Information More Valuable? An Answer With Convex Analysis

Michel De Lara[1]

[1]CERMICS, CNRS, ENPC, Institut Polytechnique de Paris, Marne-la-Vallée, France
michel.delara@enpc.fr

Keywords: value of information, information more valuable, separable additive utility, Minkowski addition

In decision problems under incomplete information, actions identified with payoff vectors indexed by (finite) states of nature and beliefs (probability distributions) are naturally paired by bilinear duality. We exploit this duality to analyze the interpersonal comparison of the value of information, using concepts and tools from convex analysis. We identify a decision-maker (DM) who maximizes expected payoff either with a closed convex positively homogeneous function over beliefs (expected payoff), or with a closed convex lower set (random payoff set). Equipped with two suitable operations that we interpret economically we obtain two isomorphic diods. Then, we show that one DM values information more than another DM if and only if their random payoff sets are solution of an algebraic equation involving the Minkowski addition, and their star-differences.

References

[1] Michel De Lara. Increasing Value of Information Implies Separable Utility. Preprint, October 2025.

[2] Michel de Lara. Adding decision problem makes information more valuable. Preprint, 2026.

[3] Michel De Lara and Olivier Gossner. Payoffs-beliefs duality and the value of information. SIAM Journal on Optimization, 30(1):464 489, 2020.

[4] Mark Whitmeyer. Making information more valuable. Journal of Political Economy, 134(3), 2026.

Session 8A: Learning Dynamics and Optimization *Room: Marie Curie (Auditorium)*

A conditional gradient approach to regularized mean field optimization problems

Pierre Lavigne*

Université Côte d'Azur, CNRS, Laboratoire J.A. Dieudonné, 06108 Nice, France; Emails: pierre.lavigne@univ-cotedazur.fr

Keywords: Mean field optimization, Conditional gradient method, Flat derivatives on space of measure, Numerical analysis.

MSC2020. Primary: 65K10, 90C26, Secondary: 46G05.

We consider regularized mean field optimization problems,

$$\inf_{m\in\mathcal{P}(\Omega)}\{G(m)+\phi(m)\},$$

where the regularization ϕ is typically given by $\varepsilon>0$ times an entropy, a Fisher information term, or an L^p-norm. We develop a conditional gradient algorithm tailored to this class of problems and provide a unified convergence analysis across several regimes.

Algorithm 1: Conditional Gradient for Mean Field Optimisation
Require: $m_0 \in \mathcal{P}(\Omega)$.
1: for $k = 0, 1, 2, ...$ do
2: Solve linearized subproblem:

$$\mu_k = \arg\min_{\mu\in\mathcal{P}(\Omega)} \left\{ \int_\Omega \frac{\delta G}{\delta m}(m_k)(x)\, d\mu(x) + \phi(\mu) \right\}.$$

3: Choose step size $\theta_k \in [0, 1]$ (e.g., line search or predefined).
4: Update: $m_{k+1} = \theta_k\mu_k + (1 - \theta_k)m_k$.
5: end for

In particular, we establish convergence guarantees in the non-convex and convex settings, as well as under relative strong convexity and L-smoothness assumptions. In the latter regimes, we derive improved convergence rates, including sublinear and linear rates under appropriate assumptions. We illustrate the effectiveness of the proposed method through a range of numerical experiments, including the minimization of maximum mean discrepancy (MMD), single-layer neural network training, and multi-layer neural network training. These experiments demonstrate the applicability of the conditional gradient approach to a variety of regularized mean field optimization problems.

Session 5H: The moment hierachy and SDP *Room: Alan Turing (A1.134)*

The Moment-SOS Hierarchy: Applications to Reaction-Diffusion PDEs and Navier–Stokes Equations

Charlie Lebarbé, Didier Henrion, Milan Korda, Émilien Flayac, Michel Fournié

1Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, France,
{charlie.lebarbe,emilien.flayac,michel.fournie}@isae-supaero.fr 2LAAS-CNRS, Université de Toulouse, France,
{henrion,korda}@laas.fr 3Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic

Keywords: Moment-SOS Relaxations, PDEs, Optimal Control, Relaxation Gap

The moment-SOS (Sum-Of-Squares) hierarchy, originally introduced by Lasserre in [1], provides a powerful framework for solving nonconvex polynomial optimization problems to global optimality by constructing a sequence of convex semidefinite relaxations. When applied to nonlinear Partial Differential Equations (PDEs), this approach allows to compute rigorous bounds on a specified, possibly nonlinear, functional of their solutions [2]. Such bounds are particularly valuable in optimal control, as they can provide quantitative guarantees on the performance of a given controller relative to the global optimum. Furthermore, the numerical pseudo-moments generated by the hierarchy can be exploited to synthesize nonlinear feedback control laws [3]. However, a central challenge for nonlinear PDEs is the possible existence of a relaxation gap between the original nonlinear, nonconvex problem and its convex relaxation over measures. Such a gap may compromise the convergence guarantees of numerical schemes based on the moment-SOS hierarchy. In this talk, we will present recent theoretical results proving the absence of a relaxation gap for scalar nonlinear parabolic PDEs of the reaction-diffusion type [4]. We will also present numerical applications of the moment-SOS hierarchy to optimal control problems governed by reaction-diffusion PDEs and the Navier–Stokes equations.

References

[1] J. B. Lasserre, Global optimization with polynomials and the problem of moments, SIAM Journal on Optimization, 11(3):796–817, 2001.

[2] M. Korda, D. Henrion, J. B. Lasserre, Moments and convex optimization for analysis and control of nonlinear PDEs, Handbook of Numerical Analysis, 23:339–366, Elsevier, 2022.

[3] C. Lebarbé, É. Flayac, M. Fournié, D. Henrion, M. Korda, Optimal control of 1D semilinear heat equations with moment-SOS relaxations, 2025 European Control Conference (ECC), 731–736, IEEE, 2025.

[4] C. Lebarbé, É. Flayac, M. Fournié, D. Henrion, M. Korda, Concentration of measure-valued solutions for semilinear parabolic equations, arXiv:2605.23678, 2026.

Session 8H: Learning dynamics and stochasticity in optimization *Room: Alan Turing (A1.134)*

Path-conditioned training: a principled way to rescale ReLU neural networks

Arthur Lebeurrier[1], Titouan Vayer[2], Remi Gribonval[3]

[1]ENS de Lyon, CNRS, Inria, Université Claude Bernard Lyon 1, LIP, UMR 5668, 69342 Lyon cedex 07, France
[2]Inria, Rennes, France
[3]Inria, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP, UMR 5668, 69342 Lyon cedex 07, France.
Correspondence to: Arthur Lebeurrier <arthur.lebeurrier@ens-lyon.fr>.

Despite recent algorithmic advances, we still lack principled ways to leverage the well-documented rescaling symmetries in ReLU neural network parameters. While two properly rescaled weights implement the same function, the training dynamics can be dramatically different. To offer a fresh perspective on exploiting this phenomenon, we build on the recent path-lifting framework, which provides a compact factorization of ReLU networks. We introduce a geometrically motivated criterion to rescale neural network parameters whose minimization leads to a conditioning strategy that aligns a kernel in the path-lifting space with a chosen reference. We derive an efficient algorithm to perform this alignment. In the context of random network initialization, we analyze how the architecture and the initialization scale jointly impact the output of the proposed method. Numerical experiments illustrate its potential to speed up training.

Session 8C: Invited session: Mixed Integer Programming and Graphs for Challenging Problems *Room: André Marie Ampère (Amphi II)*

A Clustered Team Orienteering Problem with Synchronization

Aymeric Legros[1,2], Alain Hertz[2]

[1]École Polytechnique, France, aymeric.legros@polytechnique.edu
[2]École Polytechnique Montréal, alain.hertz@gerad.ca

Keywords: VRP, Synchronization, Graphs, Integer Linear Programming

In an industrial context, this talk addresses a multi-period scheduling problem involving the routing of agents required to perform localized tasks, some of which require the simultaneous presence of multiple agents. The objective is to maximize the number of completed tasks, weighted by their priority. The proposed model relies on two key assumptions: a time discretization that facilitates the handling of synchronization constraints, and a spatial aggregation of tasks by city, within which travel distances are neglected. The aggregation of agents under the assumption of homogeneity is justified. A heuristical aggregation of tasks, leading to a formulation whose complexity is independent of the number of optional tasks, is presented.

Session 15D: Optimisation challenges in modern mobility *Room: Lise Meitner (A1.116)*

Linear Analytical Model for Static Traffic Assignment: Explicit Formulation and Optimization-Oriented Applications

Baptiste Leroux[1], Antonio Sciarretta[1], Milos Balac[2]

[1]IFP Energies Nouvelles, France, baptiste.leroux@ifpen.fr, antonio.sciarretta@ifpen.fr
[2]ETH Zurich, Switzerland, milos.balac@csfm.ethz.ch

Keywords: Static traffic assignment; Linear analytical model; Traffic equilibrium; Analytical approximation; Backcasting; Transport optimization.

This paper introduces an analytical approximation of the static traffic assignment (STA) problem that provides an explicit mapping from network characteristics and origin destination (OD) demand to link flows and travel times. The approach combines two key elements: (i) a projection of path flows onto a lower-dimensional set of auxiliary flow variables, and (ii) a local linearisation of the link performance functions around a converged traffic assignment. Under a fixed path support, these ingredients lead to a closed-form affine representation of link flows and OD travel times, together with their analytical sensitivities with respect to network and demand variables, avoiding the repeated iterative assignment procedures required by conventional STA solvers. The analytical model is validated against the Method of Successive Averages (MSA) on the Sioux Falls benchmark and on the real-scale road network of Lyon, France, comprising 8,140 links, 4,128 nodes and up to 33,258 OD pairs. The results show a high agreement with the numerical reference, with an R2 of approximately 0.99 for OD travel-time prediction. Beyond accelerating the assignment itself, the main interest of the proposed framework lies in its integration into larger optimization problems. By replacing the iterative lower-level traffic assignment with an explicit analytical mapping, the model provides a tractable surrogate for bilevel and nested optimization problems in which traffic equilibrium is repeatedly evaluated. This opens the way to gradient-based formulations of inverse and network-design problems, including OD demand calibration, network parameter estimation and graph reduction. The proposed framework therefore provides a computationally efficient bridge between traffic assignment and large-scale mathematical optimization, with particular relevance to backcasting-oriented transport planning and other decision-making problems involving repeated traffic assignment evaluations.

Session 12D: Optimal Transport & Distributional Robustness *Room: Lise Meitner (A1.116)*

Convergence of empirical subgradients for optimal transport-based objectives

Tam Le[1]

[1]Université Paris Cité, LPSM, France, tamle@lpsm.paris

Keywords: Optimal transport, empirical subgradients, subdifferential convergence, nonsmooth optimization, sliced Wasserstein costs

Optimal transport is widely used in machine learning for learning distributions, enforcing distributional constraints [3], and modeling risk aversion [2]. In one dimension, transport costs can be computed in closed form by sorting values, enabling efficient differentiation and making them practical components of learning pipelines [4, 5]. These costs are usually approximated through sampling. We study parameterized functions defined by transport costs and show that, as the sample size grows, the graphs of their subdifferentials converge [1]. In particular, this result guarantees that standard subgradient methods consistently approach population-level critical points. We illustrate our results in several settings, including risk-averse optimization, fairness, and sliced Wasserstein costs.

References

[1] T. Le, Convergence of empirical subgradients for optimal transport-based objectives, arXiv preprint arXiv:2605.28134, 2026.

[2] R. Mehta, V. Roulet, K. Pillutla, L. Liu, and Z. Harchaoui, Stochastic optimization for spectral risk measures, in Proceedings of the International Conference on Artificial Intelligence and Statistics, pp. 10112–10159, PMLR, 2023.

[3] P. Gordaliza, E. del Barrio, G. Fabrice, and J.-M. Loubes, Obtaining fairness using optimal transport theory, in Proceedings of the International Conference on Machine Learning, pp. 2357–2365, PMLR, 2019.

[4] A. Lobashev, M. Larchenko, and D. Guskov, Color conditional generation with sliced Wasserstein guidance, in Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems, 2025.

[5] L. Chapel, R. Tavenard, and S. Vaiter, Differentiable generalized sliced Wasserstein plans, in Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems, 2025.

[6] J. Rabin, G. Peyré, J. Delon, and M. Bernot, Wasserstein barycenter and its application to texture mixing, in Scale Space and Variational Methods in Computer Vision, Springer, pp. 435–446, 2011.

Session 5G: Invited Session : Optimal trading in intraday electricity markets *Room: Enrico Fermi (A1.133)*

Statistics and optimization in intraday electricity markets

Tanguy Lindegaard

Natixis, France

Keywords: Battery storage systems, Intraday trading, Markov chain, Order book

Optimal management of battery storage in the continuous intraday power market requires accounting for future order book uncertainties, rather than only reacting to their current state as myopic policies do. We propose a simple probabilistic model of the order book before gate closure to compute dynamic trading policies. We first reconstruct limit order books from order- level data for the French and German markets. We then document several stylized facts of the order books. Building on these facts, we calibrate a Markov chain on the order book state before gate closure. This chain serves as the stochastic process in an optimisation model solved by Stochastic Dual Dynamic Programming (SDDP), from which we derive a dynamic trading policy. This work connects statistical analysis of order books with stochastic control for battery storage in European intraday power markets.

Session 9B: Quantum algorithms *Room: Loggia (A1 - first floor)*

A Complex-Valued Continuous-Variable Quantum Approximation Optimization Algorithm (CCV-QAOA)

Raneem Madani, Abdel Lisser, Zeno Toffano

L2S, Université Paris-Saclay / CentraleSupélec, France

Keywords: Continuous-variable quantum computing, Quantum Approximate Optimization Algorithm, complex-valued optimization, variational quantum eigensolvers

Two intertwined challenges increasingly arise in modern optimization: the growing complexity and dimensionality of optimization problems, which place increasing demands on classical computational resources, and the prevalence of complex-valued decision variables in applications such as signal processing, communications, control, and quantum systems. Quantum computing offers alternative representations and search mechanisms for addressing such problems, while continuous-variable (CV) quantum systems provide a natural setting for continuous optimization through their infinite-dimensional Hilbert spaces.

In this work, we propose the Complex Continuous-Variable Quantum Approximate Optimization Algorithm (CCV-QAOA), a variational quantum framework for optimizing directly over complex-valued decision variables while also accommodating real-valued problems. The approach exploits the natural correspondence between the real and imaginary components of complex variables and the position–momentum quadratures of CV systems, allowing each complex variable to be represented within a single qumode and reducing the number of required qumodes. We evaluate CCV-QAOA on convex quadratic problems, constrained quadratic programs, the non-convex Styblinski–Tang benchmark, and complex quartic landscapes. We further investigate the effects of circuit depth and cutoff dimension on algorithmic performance.

Session 8I: Decision making under uncertainty *Room: Sophie Germain (A1.139)*

Robust stochastic optimization with risk-averse constraints via variance penalization and operator splitting.

Paul Malisani, Adrien Spagnol, Fabrice Guillemin

1IFP Energies nouvelles, France, firstname.surname@ifpen.fr

Keywords: Stochastic optimization, risk-averse constraints, robust optimization, optimal control, splitting operators.

This presentation addresses robust optimization problems with risk-averse and almost-sure constraints and makes several key contributions. First, we introduce a unified optimization framework that jointly integrates almost-sure constraints, risk-averse constraints, and variance penalization of decision variables. The variance penalization mechanism, inspired by machine learning to mitigate overfitting, is designed to improve the out-of-sample robustness of stochastic optimization solutions [1, 2]. Second, using operator-splitting techniques, we derive a dedicated algorithm for this framework: the Risk-Averse Variance-Penalized Progressive Hedging Algorithm (RAVP-PHA). Each iteration follows the same decomposition structure as the original Progressive Hedging Algorithm: one part corresponds to the parallelization of deterministic optimization subproblems, while the other part consists of a closed-form projection step. Third, we instantiate this framework and its solving algorithm in the context of stochastic optimal control problems with linear dynamics, ensuring convexity, and with both pure-state and mixed constraints, providing a complete solution methodology. Finally, the effectiveness of the proposed modeling framework and solving algorithm is demonstrated on an energy management application involving a battery system (EMS). Using real consumption and production data from a commercial facility, we show that combining risk-averse constraints with variance penalization significantly improves performance in terms of cost reduction and constraint satisfaction, outperforming standard MPC strategies. We show that most of the improvement over MPC is due to variance penalization.

References

[1] P. Malisani, A. Spagnol, V. Smis-Michel, Robust stochastic optimal control via variance penalization: application to Energy Management Systems, Computational Optimization and Applications, 2026.

[2] D. Mimouni, J. Zhu, W. de Oliveira, P. Malisani, A Comparative Study of Multi-Stage Stochastic Optimization Approaches for an Energy Management System, submitted to IEEE Transactions on Sustainable Energy, 2026.

Session 5E: Handling Large scale Stochastic Programs *Room: Antoine Lavoisier (A1.122)*

Decomposition of Convex Large Scale Stochastic Optimization Problems

Paul Malisani, Laurent Pfeiffer, François Caffier

[1]IFP Energies nouvelles, francois.caffier@ifpen.fr, paul.malisani@ifpen.fr
[2]Inria Saclay, Laboratoire des Signaux et Systèmes, CentraleSupélec
[3]ENSTA, laurent.pfeiffer@ensta.fr

Keywords: Stochastic Optimization, Large Scale Optimization, Multi-agent Optimization, Relaxation, Decentralization, Douglas-Rachford Splitting, Stochastic Optimal Control.

This work addresses the numerical resolution of an aggregative stochastic optimization problem in which agents are subject to both common and individual sources of uncertainty. It is motivated by the practical challenge of optimal, coordinated management of a large number of energy dynamical systems under uncertainty. We propose an extension of the work of [2] to a common noise framework, as well as a generalization to problems incorporating risk measures and coupling constraints across agents. We also propose a relaxation of our original problem, recovering the same relaxation gap result as the one established by the authors. Then, motivated by our practical problems, we restrict our attention to problems with noise having a finite number of scenarios and a scenario tree representation. Generalizing Rockafellar's expansion lemma [1, Lemma 1], we obtain an equivalent reformulation of the relaxed problem. It is particularly well-suited to numerical resolution via the Douglas–Rachford Splitting (DRS) algorithm, a well-known method for minimizing the sum of two convex functions. Indeed, the first part of the objective becomes separable with respect to both the agents and scenarios, and the second one admits a numerically tractable proximal operator. In particular, at each iteration of the algorithm, independent and parallelizable subproblems are solved. Therefore, the combination of the relaxation and the resulting decomposition allows us to control the complexity of our aggregative problem of interest as the number of agents and scenarios grows.

References

[1] R. T. Rockafellar, Progressive decoupling of linkages in optimization and variational inequalities with elicitable convexity or monotonicity, Set-Valued and Variational Analysis, 27(4):863–893, 2019.

[2] A. Seguret, C. Alasseur, J. F. Bonnans, A. De Paola, N. Oudjane, and V. Trovato, Decomposition of convex high dimensional aggregative stochastic control problems, Applied Mathematics & Optimization, 88(1):8, 2023.

Session 8G: Optimizing under network structure *Room: Enrico Fermi (A1.133)*

Data-Driven Security Assessment of Distribution Grids with Weibull-Based Synthetic Data Generation

Alban Marie[1], Juan J. Cuenca[2], Emanuel Aldea[3], Guy Camilleri[4], Anne Blavette[1]

[1]IETR - UMR CNRS 6164 Univ. Rennes, ENS Rennes, France, firstname.lastname@ens-rennes.fr
[2]IETR - UMR CNRS 6164 CentraleSupélec Rennes, France, juan.cuenca@centralesupelec.fr
[3]SATIE CNRS UMR 8029 Univ. Paris Saclay Gif-sur-Yvette, France, emanuel.aldea@universite-paris-saclay.fr
[4]IRIT-SMAC Univ. de Toulouse, CNRS, Toulouse INP, UT3 Toulouse, France, guy.camilleri@irit.fr

Keywords: distribution grid, data-driven security assessment, machine learning, synthetic data generation, Weibull distribution

Data-driven surrogate models enable faster security assessment of power grids compared to repeated steady-state power-flow calculations. Training such surrogate models poses a challenge of building representative labelled training data, because historical measurements are scarce due to privacy constraints. Also, constraint-violating operating conditions are rare in historical measurements, leaving few samples near the classification boundary; this imbalance in training data limits the ability of historical data alone to train a classifier for rare violation events. Synthetic-data generation strategies have therefore been proposed [2], while deep active learning can reduce the number of power-flow-labelled points by selecting informative samples [1]. Here, we propose a novel training data generation strategy that samples active power independently for each load from a truncated Weibull distribution. Its parameters control the frequency of higher-consumption operating points and are tuned to obtain a balanced proportion of constraint-satisfying and constraint-violating samples. Unlike iterative power-flow-based generation [2], the procedure is independent of the evaluation oracle and can therefore be combined with an active-learning-based labelling process [1]. The method is evaluated against uniform and iterative sampling on low- and medium-voltage test cases. Testing uses realistic residential load profiles. The proposed Weibull strategy achieves higher F1-scores across the reported cases with substantially fewer labelled points. In particular, surrogate models achieve high recall; this is desirable for security assessment because false alarms are preferable to missed constraint violations. Once trained, the ML surrogate evaluates operating points several orders of magnitude faster than conventional power-flow assessment.

References

[1] A. Marie, J. J. Cuenca, E. Aldea, G. Camilleri, and A. Blavette, Deep Active Learningdriven Acceleration of Smart Grids Security Assessment, in 2025 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), pp. 1–5, La Valletta, Malta, 2025, doi:10.1109/ISGTEurope64741.2025.11304768.

[2] J. J. Cuenca, E. Aldea, E. Le Guern-Dall'o, R. Feraud, G. Camilleri, and A. Blavette, Training Data Generation Strategies for Data-driven Security Assessment of Low Voltage Smart Grids, in 2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE), pp. 1–5, Dubrovnik, Croatia, 2024, doi:10.1109/ISGTEUROPE62998.2024.10863199.

Session 15G: Data Science, Learning and the intersection with Optimization *Room: Enrico Fermi (A1.133)*

On ℓ_0-Regularized Sparse Multiple Kernel Learning for Feature Selection

Rui Ma[1], Rodolphe Le Riche[1]

[1]LIMOS (UMR CNRS 6158), EMSE and UCA, France,
rui.ma@doctorant.uca.fr, rodolphe.le_riche@uca.fr

Keywords: Sparse Multiple Kernel Learning, Mixed-Integer Programming, $\ell 0$ norm

We study the direct use of $\ell 0$ regularization in sparse multiple kernel learning (SMKL) problems. We introduce an explicit $\ell 0$ norm on the kernel combination coefficients, leading to min–max mixed-integer optimization formulations in which kernel selection is handled in the outer problem and a standard SVM is solved in the inner problem; we solve these formulations to within a prescribed tolerance of global optimality. For feature-wise rank-one linear base kernels, we show that the number of selected kernels provides an explicit upper bound on the rank of the combined kernel. Under this construction, the $\ell 0$ constraint directly controls the number of selected kernels and admits an interpretation in terms of feature selection. Furthermore, to disentangle the sparsity-inducing and kernel-scale-controlling effects of $\ell 1$ regularization in standard MKL, we introduce an $\ell 0$–ℓp MKL framework that uses explicit $\ell 0$ regularization to control the selected support and continuous ℓp regularization to control the magnitude and scale of the kernel weights. We compare our $\ell 0$-MKL formulations with conventional $\ell 1$-based MKL [1, 2] and a recent heuristic approach [3] for $\ell 0$-based MKL. Extensive experiments on synthetic and UCI benchmark datasets demonstrate that the proposed $\ell 0$-regularized SMKL approach selects fewer kernels than conventional $\ell 1$-MKL while maintaining comparable test accuracy and generally achieving higher predictive accuracy than the alternating best-response heuristic. These results complement previous observations in MKL where less sparse solutions are often observed to perform better.

References

[1] Sören Sonnenburg, Gunnar Rätsch, Christin Schäfer, and Bernhard Schölkopf. Large scale multiple kernel learning. The Journal of Machine Learning Research, 7:1531–1565, 2006.

[2] Alain Rakotomamonjy, Francis Bach, Stéphane Canu, and Yves Grandvalet. SimpleMKL. Journal of Machine Learning Research, 9:2491–2521, 2008.

[3] Dimitris Bertsimas, Caio de Próspero Iglesias, and Nicholas A. G. Johnson. Sparse multiple kernel learning: Alternating best response and semidefinite relaxations. Transactions on Machine Learning Research, 2025. ISSN 2835-8856. URL https://openreview.net/forum? id=Y5icwFwkyh.

Session 5A: Invited session: Nonconvex landscapes and algorithms in data science *Room: Marie Curie (Auditorium)*

Attention-based PCA

Rodrigo Maulen-Soto[♠], Claire Boyer[♢♣]

[♠]Sorbonne Université, CNRS, Laboratoire de Probabilités, Statistique et Modélisation, 75005 Paris, France
[♢]Université Paris-Saclay, CNRS, Inria, Laboratoire de mathématiques d'Orsay, 91405 Orsay, France
[♣]Institut Universitaire de France

June 19, 2026

We study attention mechanisms through the lens of a canonical unsupervised problem: principal component analysis (PCA). We show that, when trained on Gaussian data, both softmax and linear attention layers learn parameters that align with the principal eigenvectors of the covariance matrix, thereby establishing a direct and explicit connection with PCA. Our analysis covers both finite and infinite prompt regimes. In the infinite-prompt limit, we prove convergence to globally optimal solutions aligned with the leading spectral direction, while in the finiteprompt setting we show that the same behavior emerges up to sampling effects. We further extend the analysis to an in-context setting with spiked Wishart covariances, where attention successfully recovers the underlying signal direction. These results demonstrate that attention inherently performs PCA-like computations under unsupervised objectives, providing a theoretical foundation for its representation-learning capabilities.

Session 8D: Invited session: Recent Advances in Systematic and Computer-Aided Methods in Optimization part 2
Room: Lise Meitner (A1.116)

opt-syn: Analysis and Synthesis of First-Order Algorithms

Jared Miller[1], Manuel Zobel[1], Fabian Jakob[2], Andrea Iannelli[2], Carsten Scherer[1]

[1]Chair of Mathematical Systems Theory, University of Stuttgart, Germany
{jared.miller,carsten.scherer}@imng.uni-stuttgart.de
[2]Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany
{fabian.jakob,andrea.iannelli}@imng.uni-stuttgart.de

Keywords: Optimization Algorithms, Analysis, Design, Robust Control, Software

opt-syn (https://opt-syn.github.io/) is a MATLAB package for the automated analysis and synthesis of first-order algorithms. Optimization and fixed-point algorithms can be interpreted through the lens of dynamical systems as the interconnection of linear systems and operator nonlinearities [1]. The convergence of well-posed optimization algorithms can then be connected to disturbance rejection/output regulation properties of these nonlinear systems [2, 3]. This dynamical systems formulation allows for the analysis and synthesis of optimization algorithms by solving robust control problems [4]. Analysis bounds performance measures such as worst-case linear convergence rates, and gains for robust or stochastic error amplification. Synthesis tries to design optimization algorithm satisfying these desired performance constraints. Analysis and Synthesis may be performed for algorithms arising in dynamic environments, including cases with constant or time-varying delays, channel memory, and cross-talk. The Analysis and Synthesis problems are posed using the frameworks of Integral Quadratic Constraints [5, 6, 7], and are solved through Linear Matrix Inequality methods [8].

References

[1] J. Wang and N. Elia, "A Control Perspective for Centralized and Distributed Convex Optimization," in 2011 Proc. IEEE Conf. Decis. Control. IEEE, 2011, pp. 3800–3805.

[2] J. Miller, C. Scherer, F. Jakob, and A. Iannelli, "Structure, analysis, and synthesis of first-order algorithms," arXiv preprint arXiv:2603.24795, 2026.

[3] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," Automatica, vol. 12, no. 5, pp. 457–465, 1976.

[4] C. W. Scherer and C. Ebenbauer, "A Tutorial on Convex Design of Optimization Algorithms by Integral Quadratic Constraints," Annual Review of Control, Robotics, and Autonomous Systems, vol. 9, 2025.

[5] A. Megretski and A. Rantzer, "System analysis via integral quadratic constraints," IEEE Transactions on Automatic Control, vol. 42, no. 6, pp. 819–830, 1997.

[6] L. Lessard, B. Recht, and A. Packard, "Analysis and Design of Optimization Algorithms via Integral Quadratic Constraints," SIAM Journ. Opt., vol. 26, no. 1, pp. 57–95, 2016.

[7] J. Veenman, C. W. Scherer, and H. Köroğlu, "Robust stability and performance analysis based on integral quadratic constraints," European Journal of Control, vol. 31, pp. 1–32, 2016.

[8] C. Scherer and S. Weiland, "Linear Matrix Inequalities in Control," Lecture notes, dutch institute for systems and control, Delft, the Netherlands, vol. 3, no. 2, pp. 62–74, 2000.

Session 15J: Operations Research for Food and Healthcare *Room: Emmy Noether (A1.140)*

Optimizing Hospital Procurement Decisions: The Case of a University Hospital

Cecília Mota, Lara Pontes, Luciano Costa, Anand Subramanian

1Universidade Federal da Paraíba, Brazil, anaceciliabezerra10@gmail.com, luciano.costa@academico.ufpb.br, anand@ci.ufpb.br 2Massachusetts Institute of Technology, USA, lpontes@mit.edu

Keywords: Hospital inventory management. Operations research in healthcare. Integer programming.

Hospital inventory management is challenged by budget constraints, demand variability, and the large variety of items. These items also differ in their levels of criticality, requiring management approaches that account for their impact on healthcare services, such as the VED (Vital, Essential, and Desirable) classification [1]. Moreover, stock-outs may compromise the continuity of care and patient safety [2]. In Brazilian public hospitals, these challenges are further compounded by bureaucratic procurement processes, in which the same item may be associated with different price registration agreements, each with distinct suppliers, prices, available quantities, and priority levels. In this problem, the purchasing decision involves determining how much of each product to acquire and from which of the available price registration agreements, while considering the desired inventory level based on product demand, the criticality of each item, and the priorities among available agreements. This work proposes two mathematical programming formulations with multiple objectives to support purchasing decisions in a university hospital, integrating operational and legal aspects of the procurement process with product prioritization according to item criticality. Both models determine what and how much to acquire while considering budget limitations, agreement priorities, and penalties for understocking. The first model considers a single purchasing period, oriented toward operational decisions and emulating current purchasing practice, while the second adopts a multi-period formulation that supports more tactical and strategic planning over a time window. The approaches were evaluated using real hospital instances constructed from historical data on inventory levels, consumption, and active price registration agreements. The results indicate that the proposed models can generate replenishment decisions that maintain inventory coverage within the desired range while respecting budget and procurement constraints, with more stable inventory levels than those observed historically.

References

[1] R. Gupta, K. K. Gupta, B. R. Jain, R. K. Garg„ ABC and VED Analysis in Medical Stores Inventory Control, Medical Journal Armed Forces India, 63(4):325–327, 2007.

[2] R. L. Mendoza, Materials management under conditions of uncertainty: Does it matter much in healthcare?, International Journal of Healthcare Management, 18(3):641–652, 2025.

Session 15G: Data Science, Learning and the intersection with Optimization *Room: Enrico Fermi (A1.133)*

Multimodal Optimization of K-means Clustering

Ravil Mussabayev, Rustam Mussabayev

1AI Research Lab, Satbayev University, Almaty, Kazakhstan r.mussabayev@satbayev.university, ru.mussabayev@satbayev.university

Keywords: K-means clustering, big data clustering, multimodal optimization, variable landscape search.

The minimum sum-of-squares clustering (MSSC) problem, also known as K-means clustering, is a fundamental NP-hard problem in operations research [1]. Given a dataset $X = \{x_1, \dots, x_m\} \subset \mathbb{R}^n$, it seeks k centroids $C = (c_1, \dots, c_k)$ that minimize

$$\min_{C \in \mathbb{R}^{n \times k}} f(C, X) = \sum_{i=1}^{m} \min_{j=1,\dots,k} \|x_i - c_j\|^2.$$

Its nonconvex landscape contains many local minima, while big datasets make extensive exploration computationally expensive. We propose a multimodal perspective on large-scale MSSC optimization [2]. Here, an optimization modality is a distinct way of influencing the search: (1) formulating the objective, (2) selecting and representing the data, or (3) exploring the solution landscape. Combining these modalities provides complementary mechanisms for escaping local minima. Our studies show that combining data selection and landscape exploration improves solution quality and computational efficiency over corresponding unimodal methods. Big-means [1] and Big-means++ [3] repeatedly optimize MSSC on small random subsets of the dataset. Preliminary, unpublished results indicate that adding a collective intelligence method for more extensive landscape exploration yields further performance gains. A promising direction is to integrate objective reformulations, such as adding weighted regularization terms, changing the coordinate system, or even switching to alternative clustering models, with data selection and landscape exploration modalities.

References

[1] R. Mussabayev, N. Mladenovic, B. Jarboui, R. Mussabayev, How to Use K-means for Big Data Clustering?, Pattern Recognition, 137 (2023), 109269.

[2] R. Mussabayev, R. Mussabayev, Variable Landscape Search: A Novel Metaheuristic Paradigm for Unlocking Hidden Dimensions in Global Optimization, Mathematics, 13(22), 3667, 2025.

[3] R. Mussabayev, R. Mussabayev, Z. Yerdaliyeva, K. Nursultan, Data-Native Global Optimization for Big Data K-means Clustering, 2026, arXiv:2607.15835.

Session 12K: Multiobjective Optimization and Evolutionary Computation Part I *Room: Loggia (A1 - first floor)*

Discovering Interpretable Multi-Parameter Control Policies for Evolutionary Algorithms Using Deep Reinforcement Learning

Tai Nguyen, Phong Le, Carola Doerr, Nguyen Dang

1University of St Andrews, St Andrews, United Kingdom 2Sorbonne University, CNRS, LIP6, Paris, France

Keywords: parameter control, dynamic algorithm configuration, reinforcement learning

Parameter control aims at dynamically adapting the behavior of an optimization algorithm to the different stages of the search [1]. While dynamic parameter choices can provide substantial benefits over static configurations, deriving effective multi-parameter control policies remains challenging, especially when the policy must both perform well empirically and remain interpretable enough to support theoretical analysis. In this talk, we present an investigation of deep reinforcement learning (deep RL) as a potential tool for discovering such policies for evolutionary algorithms. We consider the $(1+(\lambda, \lambda))$-Genetic Algorithm on the OneMax benchmark, a well-studied setting in which dynamic parameter control is known to yield a provable asymptotic advantage [2]. We formulate the control of its four main parameters: the mutation and crossover population sizes, the mutation-rate coefficient, and the crossover-bias coefficient, as a Dynamic Algorithm Configuration (DAC) problem [3]. Using Double Deep Q-Networks (DDQN), we show that deep RL can discover highly effective multi-parameter policies, despite the non-trivial structure of the resulting decision problem. A simple policy derived from the learned behavior consistently outperforms previously proposed theoretical and algorithmically tuned policies on problem sizes up to $n = 40\,000$. More importantly, we address a key limitation of deep-RL-based control: although neural policies can be highly effective, they are difficult to interpret and analyze theoretically. We therefore develop a two-stage distillation procedure that transforms learned policies into explicit symbolic control rules. The resulting white-box policies expose interpretable relationships between the algorithm's parameters and the optimization state while retaining strong empirical performance. We further automatically fine-tune the symbolic policy using SMAC3 [4], obtaining additional improvements over the initial symbolic representation. Together, these results demonstrate a data-driven route from black-box policy discovery to interpretable algorithm design. Beyond the considered benchmark, this approach suggests that deep RL can serve not only to optimize evolutionary algorithms, but also to generate concrete mathematical hypotheses for the analysis of adaptive optimization methods.

References

[1] A. E. Eiben, Z. Michalewicz, M. Schoenauer, and J. E. Smith, "Parameter control in evolutionary algorithms," in Parameter Setting in Evolutionary Algorithms, 2007, pp. 19–46.

[2] B. Doerr and C. Doerr, "Optimal static and self-adjusting parameter choices for the $(1 + (\lambda, \lambda))$ genetic algorithm," Algorithmica, 80, 1658–1709, 2018.

[3] A. Biedenkapp, N. Dang, M. S. Krejca, F. Hutter, and C. Doerr, "Theory-inspired parameter control benchmarks for dynamic algorithm configuration," in Proceedings of GECCO, 2022, pp. 766–775.

[4] M. Lindauer, K. Eggensperger, M. Feurer, A. Biedenkapp, D. Deng, C. Benjamins, T. Ruhkopf, R. Sass, and F. Hutter, "SMAC3: A versatile Bayesian optimization package for hyperparameter optimization," Journal of Machine Learning Research, 23, 1–9, 2022.

Session 12K: Multiobjective Optimization and Evolutionary Computation Part I *Room: Loggia (A1 - first floor)*

Reinforcement Learning for Solving Multimodal Functions with CMA-ES

Marc Oedingen, Oskar Girardin, Nikolaus Hansen, Alexandre Chotard

1TH Köln, Institute for Data Science, Engineering, and Analytics, Germany, marc.oedigen@th-koeln.de 2Inria, CMAP École polytechnique, France, firstname.lastname@inria.fr 3Université du Littoral Côte d'Opale, France, alexandre.chotard@univ-littoral.fr

Keywords: CMA-ES, Dynamic Parameter Control, Reinforcement Learning, Black-box Optimization

In this paper, we investigate dynamic meta-parameter control for the covariance matrix adaptation evolution strategy (CMA-ES) through deep reinforcement learning (DRL). We formalize the online adaptation of meta-parameters as a Markov decision process (MDP), where each iteration corresponds to one decision step. Based on this formalization, we integrate a controller designed with proximal policy optimization (PPO) into CMA-ES that adapts the learning rates of the mean and the damping parameter of cumulative step-size adaptation (CSA) during the optimization run. Importantly, the controller preserves all known invariance properties of CMA-ES. We train the PPO model only on the Rastrigin function. The learned controller enables CMA-ES with default population size to solve a variety of multimodal functions: on the Rastrigin function, the PPO-CMA-ES is four times faster than IPOP-CMA-ES, and it is also highly competitive on a variety of other multimodal functions, while maintaining reasonable performance on unimodal functions.

References

[1] Marc Oedingen, Oskar Girardin, Nikolaus Hansen, Alexandre Chotard. Reinforcement Learning for Solving Multimodal Functions with CMA-ES. PPSN 2026 - 19th International Conference on Parallel Problem Solving from Nature, University of Trento, Aug 2026, Trento, Italy. pp.136-151 10.1007/978-3-032-36217-9_9. hal-05744963

Session 13C: PDE-constrained optimization *Room: André Marie Ampère (Amphi II)*

Rank-Adaptive Tensorized Neural Operators for PDE-Constrained Optimization

Younes Brahim Oumedjber[1], Akram Boukabache[2], Nadjette Debbouche[2], Adel Ouannas[3]

[1]Laboratory of Materials Science for Applications and the Environment (LaSMAE), Department of Sciences, Teacher Education College of Sétif–Messaoud Zeghar, El Eulma, Sétif, Algeria, younes-brahim.oumedjber@ens-setif.dz

[2]Department of Sciences, Teacher Education College of Sétif–Messaoud Zeghar, El Eulma, Sétif, Algeria

[3]Department of Mathematics, Faculty of Science, Islamic University of Madinah, Madinah, Saudi Arabia

Keywords: PDE-constrained optimization, neural operators, tensor trains, scientific machine learning, low-rank optimization

PDE-constrained optimization requires repeated solutions of state and adjoint equations and can become prohibitive in many-query, high-dimensional, or uncertainty-aware settings. Neural operators [1] offer an attractive amortized surrogate, but dense spectral kernels remain over-parameterized, while fixed low-rank compression can degrade the derivatives that drive optimization [3]. We propose a rank-adaptive tensorized neural operator designed specifically for surrogate-based PDE-constrained optimization. For a parameterized problem

$$\min_{z \in Z} J(u, z) \qquad \text{subject to} \qquad A(m, z)u = f.$$

the solution map (m, z) 7→u is learned by a Fourier-type neural operator whose high-dimensional spectral weights are represented in Tensor-Train (TT) form [2]. Instead of fixing the TT ranks a priori, we treat them as discrete complexity variables and adapt them using joint state, derivative, and compression indicators. The training objective combines solution error, a physics residual, gradient-consistency information relevant to the reduced optimization problem, and an explicit TT-complexity penalty. This produces a coupled accuracy–complexity mechanism in which ranks are increased only where additional expressive power improves the optimization-relevant validation criterion and are truncated when factor energy becomes negligible. The surrogate is embedded in a gradient-based optimization loop and evaluated on parametric elliptic/Darcy-type control problems. The numerical study compares a full Fourier neural operator, fixed-rank TT surrogates, and the proposed adaptive model using state relative L2

error, reduced-gradient error, objective optimality gap, PDE feasibility residual, parameter count, memory footprint, and wall-clock cost. Additional tests under parameter distribution shift assess whether adaptive low-rank structure acts as a regularizer rather than merely a compression device. The framework connects PDE-constrained optimization, numerical analysis, low-rank tensor approximation, and scientific machine learning, with the aim of obtaining compact surrogates whose derivatives remain reliable enough for optimization.

References

[1] Z. Li et al., Fourier Neural Operator for Parametric Partial Differential Equations, ICLR, 2021. [2] I. V. Oseledets, Tensor-Train Decomposition, SIAM J. Sci. Comput. 33(5):2295–2317, 2011. [3] Z. Cheng et al., Accelerating PDE-Constrained Optimization by the Derivative of Neural Operators, ICML, 2025.

Session 8F: Solving games and Equilibrium problems *Room: Ada Lovelace (A1.128)*

Interior-point methods for large-scale equilibrium problems

François Pacaud, Benoît Legat, Quentin Lété, Quentin Jacquet

[1]Centre Automatique et Systèmes, Mines Paris-PSL, France; [2]Université catholique de Louvain, Belgique; [3]Osiris, EDF R&D, France

Keywords: Equilibrium problems; Bilevel programming; Energy markets; Interior-point methods

Computing the market equilibrium involves solving a set of variational inequalities, which is nontrivial to do accurately. The problem is here formulated as a large-scale Equilibrium Problem with Equilibrium Constraints (EPEC). We propose a solution based on nonlinear programming, where the EPEC is interpreted as a particular instance of a Mathematical Program with Com- plementarity Constraints (MPCC). We develop a package to formulate the market equilibrium directly in the JuMP modeling language and implement automatic procedures to reformulate the problem as an MPCC. In particular, we compare two specific reformulations for the result- ing MPCC: (i) the Scholtes relaxation, which reformulates the MPCC as a (nearly degenerate) nonlinear program, allowing it to be solved with a nonlinear interior-point solver such as Ipopt; (ii) a reformulation of the complementarity constraints using SOS1 constraints, as is standard in the literature. We test both approaches numerically on several practical EPEC instances: (i) a tariff-setting problem given by EDF, (ii) investment under risk in incomplete markets, (iii) wholesale market equilibrium with market power, (iv) wholesale market equilibrium with day-ahead/real-time arbitrage. We assess the strengths and weaknesses of each method.

Session 8K: Invited session : Quantum Applications in Operations Research part 2 *Room: Loggia (A1 - first floor)*

Quantum Computing for Integer Quadratic Programming

Dimitri Papadimitriou

Universite Libre de Bruxelles, Belgium

A principled way to solve NP-hard combinatorial optimization problems by means of a quantum computing algorithm is to reformulate them as quadratic unconstrained binary optimization (QUBO) problems. Formally, a QUBO is defined by

$$\min_{x\in\{0,1\}^n} x^\top Q x + c^\top x \quad \Longleftrightarrow \quad \min_{x\in\{0,1\}^n} x^\top (Q + \mathrm{diag}(c))x = x^\top Q' x.$$

Here $Q \in \mathbb{R}^{n\times n}$ is symmetric and encodes pairwise interactions, while $c \in \mathbb{R}^n$. Since $x_i^2 = x_i$ for binary variables, the linear coefficients can be absorbed into the diagonal of Q'. The constrained binary problem can also be written in Ising variables $s \in \{-1, 1\}^n$ through $s = 2x - \mathbf{1}$:

$$\min_{s\in\{-1,1\}^n} s^\top J s + h^\top s.$$

This transformation allows QUBOs to be attacked by quantum annealing or variational algorithms such as QAOA [1]. Although the QUBO and Ising representations are isomorphic, linear terms cannot be absorbed into the diagonal of J: because $s_i^2 = 1$, one has $s^\top \mathrm{diag}(h) s = \sum_i h_i$, a constant. Hence the term $h^\top s$ must remain explicit; misrepresenting it changes energy differences and may change the ground state.

One may distinguish problems with a natural QUBO form from constrained binary problems for which constraints must be moved to the objective using penalty terms. The predominant approach for linearly constrained problems is to penalize the quadratic violation of the constraints, often after introducing binary-encoded slack variables for inequalities.

We show that, for this latter class, one should not be limited to quadratic penalties: the linear violation can also be penalized. In mixed-integer and nonlinear programming this is the augmented Lagrangian method (ALM) introduced by Powell and Hestenes [2,4], alternating primal minimization of the augmented Lagrangian with dual-gradient updates of the multipliers. An augmented-Lagrangian QUBO variant was recently considered for the quadratic knapsack problem on a quantum annealer [5]. We characterize penalty terms for general classes of linearly constrained problems that preserve a QUBO form and illustrate them on max-k-cut and quadratic-knapsack variants. We also develop alternative formulations for inequality violations that mitigate the large increase in binary variables caused by standard slack encodings [3].

References

[1] E. Farhi, J. Goldstone, S. Gutmann, "A quantum approximate optimization algorithm," arXiv:1411.4028.
[2] M. R. Hestenes, "Multiplier and gradient methods," J. Optim. Theory Appl. 4 (1969), 303–320.
[3] A. Lucas, "Ising formulations of many NP problems," Frontiers in Physics 2:5 (2014).
[4] M. J. D. Powell, "A method for non-linear constraints in minimization problems," in R. Fletcher (ed.), Optimization, Academic Press, 1969, 283–298.
[5] K. Yonaga, M. J. Miyama, M. Ohzeki, "Solving inequality-constrained binary optimization problems on quantum annealer," arXiv:2012.06119.

Session 5G: Invited Session : Optimal trading in intraday electricity markets *Room: Enrico Fermi (A1.133)*

Convex optimization in intraday electricity markets

Teemu Pennanen

King's College London, UK

Keywords: Battery storage systems, Intraday trading, Dynamic programming, Convex programming

We study optimal management of Battery Energy Storage Systems in intraday power mar- kets. A popular approach is the so-called "rolling intrinsic" strategy, which optimises a determin- istic trading strategy against the available intraday price quotes. This results in deterministic optimisation problems that are particularly easy to solve, but the obtained "intrinsic value" misses the value of optionality of a battery, which may be significant in the face of price uncer- tainty. We propose stochastic optimisation models that 1) allow for dynamic trading strategies that may react to the arrival of new information on prices, etc., 2) can be solved efficiently using techniques of dynamic programming and convex optimisation, and 3) outperform the determin- istic formulations, which are obtained as special cases.

Session 12I: Optimization for Learning and Large-Scale Decision Making *Room: Sophie Germain (A1.139)*

Differentiating Minimal-Norm Solutions to Parametric Optimization Problems

Baptiste Plaquevent-Jourdain[1], Jalal Fadili[2], Antonio Silveti-Falls[1]

[1]CVN-OPIS, CentraleSupélec, France, baptiste.plaquevent-jourdain@USherbrooke.ca, tonys.falls@gmail.com
[2]GREYC CNRS, France, jalal.fadili@ensicaen.fr

Keywords: bilevel programming, implicit differentiation, conservative mappings, minimalnorm solutions, Tikhonov regularization)

Differentiating through parametric optimization problems is central to bilevel programming and meta-learning, often accomplished using approximate implicit differentiation. The implicit function theorem requires inverting a partial Jacobian of the optimality condition, which fails when there are many solutions. Nonetheless, in such cases it is possible to relax invertibility to a strictly weaker uniform range condition, under which it is shown that the minimal-norm solution mapping admits generalized derivatives by using a limiting Tikhonov regularization argument and conservative set-valued field theory [1, 2]. With additional control on the eigenvalues of the generalized Hessians, a pseudoinverse formula is justified. This is established for a class of smooth convex objectives and extended to nonsmooth composite problems. These assumptions are verified for Least-Squares, Huber regression and LASSO. The resulting extension of nonsmooth implicit differentiation to ill-posed settings is examined experimentally on data poisoning and data hypercleaning problems. A preprint is available on Arxiv, see [3].

References

[1] Jérôme Bolte and Édouard Pauwels, Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning, Mathematical Programming 188 (April 2020), 19–51.

[2] Sholom Schechtman, The gradient's limit of a definable family of functions admits a variational stratification, SIAM Journal on Optimization 36 (2026), 1075–1099

[3] Baptiste Plaquevent-Jourdain, Jalal Fadili, Antonio Silveti-Falls, Differentiating MinimalNorm Solutions to Parametric Optimization Problems, arXiv:260828899

Session 8I: Decision making under uncertainty *Room: Sophie Germain (A1.139)*

Probing-Enhanced Stochastic Programming: Structural Dominance, Stability, and Approximate Branch-and-Bound

Franco Quezada, Céline Gicquel, Safia Kedad-Sidhoum, Bernardo Pagnoncelli

1University of Santiago of Chile (USACH), Chile, franco.quezada@usach.cl 2Université Paris-Saclay, LISN, France, gicquel@lisn.fr 3Conservatoire National des Arts et Métiers, CEDRIC, France, safia.kedad_sidhoum@cnam.fr 4SKEMA Business School, Université Côte d'Azur, France, bernardo.pagnoncelli@skema.edu

Keywords: decision-dependent uncertainty, information acquisition, stochastic programming, branch-and-bound, lot-sizing

We study stochastic optimization problems in which a decision-maker can acquire partial information about uncertainty before making operational decisions. Selecting what information to collect is combinatorial, and each selection induces a different conditional stochastic optimization problem. This combination makes probing-enhanced stochastic programs particularly challenging, because many closely related subproblems must be solved repeatedly while exploring alternative information structures. We develop structural and algorithmic tools that exploit the relationships among these information structures. First, we characterize probing decisions through the scenario partitions they induce and derive dominance relations that identify redundant information sources and reduce the effective search space. Second, we establish stability and solution-reuse properties under changes in the information partition, allowing decisions computed for previously solved subproblems to be inherited or certified. Building on these results, we propose an approximate information-relaxation branch-and-bound algorithm that combines dominance-based pruning with a hierarchical evaluation of aggregated subproblems. Exact optimization is used only when inherited solutions and progressively refined bounds cannot provide the required certificate. We apply the framework to a multi-item capacitated lot-sizing problem in which demandrelated signals can be selectively observed before setup and production decisions are made. Computational experiments show that dominance rules reduce the average number of explored nodes by up to 47.6%. With a moderate approximation tolerance, only 16.2% of block subproblems require exact evaluation. The proposed method achieves runtime reductions exceeding a factor of twenty while keeping gaps below 0.83%, and it produces feasible bounds for instances with up to 3,200 scenarios that the exact baseline generally cannot handle. A complementary analysis shows that the relative operational gain from optimal probing reaches 45.13%, highlighting both the computational and managerial value of selective information acquisition.

Session 8F: Solving games and Equilibrium problems *Room: Ada Lovelace (A1.128)*

A New Solution Concept for Games Allowing for Infeasibility

Anne Reulke[1,2], Mikaël Touati[1], Rachid El-Azouzi[2]

[1]Orange Research, France, anne.reulke@orange.com, mikael.touati@orange.com
[2]Laboratoire d'Informatique d'Avignon, France, rachid.elazouzi@univ-avignon.fr

Keywords: Equilibrium, Infeasibility, Generalized Games, Penalty Method

Addressing infeasibility in non-cooperative games has become an important topic, as many problems across different applications face this issue. We propose a new solution concept for generalized games with possibly infeasible individual constraints. A solution is defined as the limit of a sequence of generalized Nash equilibria induced by games with penalty terms relaxing the individual constraints. Existence is established for a broad range of games and we provide conditions allowing to characterize a ψ-penalized solution as a strategy profile maximizing every player's utility over all her penalty minimizing strategies. A variation of Divide-the-Dollar serves as an illustrative example.

Session 12F: Routing & Mobility *Room: Ada Lovelace (A1.128)*

Robotaxi Dispatching with Endogenous En-route Times

David Rey[1]

[1]SKEMA Business School, Université Côte d'Azur, France, david.rey@skema.edu

Keywords: robotaxi dispatching, endogenous times, network flow, nonconvex optimization, spatial branch-and-bound

We consider the problem of routing robotaxis, i.e., autonomous vehicles for taxi transportation, within in a city. We assume that the city is divided into regions and that demand is organized in origin-destination markets. We model the operations of robotaxis in this network under steady-state conditions and we aim to determine optimal dispatching and repositioning strategies. In this model, robotaxis can be in one of four states: occupied (transporting passengers), en-route to pickup, repositioning or idle. We assume that vehicle-passenger matching is endogenous to the state of the system, i.e., in each region, the number of idle vehicles affects traveler's waiting time and thus service quality [1]. To model the relationship between trip demand and vehicle supply in each region, we adopt the stock-based matching function from the transit systems literature which is based on the nearest idle vehicle geometry, also known as the square-root law [2, 3]. We embed the stock-based matching function within a network flow model to represent multi-region vehicle dispatching operations. This leads to a nonconvex mathematical optimization formulation. We conduct a theoretical analysis of this mathematical optimization problem and derive structural results. We characterize the relationship between each region's idle stock and routing rates and study under which conditions are markets served. To solve this mathematical optimization problem, we develop a spatial branch-and-bound algorithm that exploits several convex relaxations. In particular, we identify a relaxation of the matching function that is second-order cone representable [4] and is binding at the optimum. This allows us to confine the nonconvexity to region-specific bilinear terms which can be handled through McCormick relaxations. We report preliminary results on problem instances based on realistic urban transportation networks that illustrate the impact of endogenous en-routes times onto system performance and service quality.

References

[1] C. F. Daganzo, Y. Ouyang, A general model of demand-responsive transportation services, Transp. Res. B, 126:213–224, 2019.

[2] R. Arnott, Taxi travel should be subsidized, J. Urban Econ., 40(3):316–333, 1996.

[3] A. Kleywegt, H. Shao, Optimizing pricing, repositioning, en-route time, and idle time in ride-hailing systems, 2021, arXiv:2111.11551.

[4] M. S. Lobo, L. Vandenberghe, S. Boyd, H. Lebret, Applications of second-order cone programming, Linear Algebra Appl., 284:193–228, 1998.

Session 5I: Invited Session:Neural Network approximations for high-dimensional PDEs Problems*Room: Sophie Germain (A1.139)*

Global Convergence of Second-Order Methods for Physics-Informed Neural Networks

Konstantin Riedl[1], Justin Sirignano[2], Konstantinos Spiliopoulos[3]

[1]University of Oxford, United Kingdom, konstantin.riedl@maths.ox.ac.uk
[2]University of Oxford, United Kingdom, justin.sirignano@maths.ox.ac.uk
[3]Boston University, USA, kspiliop@bu.edu

Keywords: physics-informed neural network, second-order optimizer, overparameterized limit

Solving high-dimensional partial differential equations (PDEs) constitutes a core task in many engineering and scientific disciplines. While classical methods such as the finite difference, finite element, or spectral element method suffer from the curse of dimensionality, physics-informed neural networks and the deep Galerkin method effectively leverage machine learning models to approximate solutions of known PDEs by minimizing residuals corresponding to the PDE operator alongside external constraints. While second-order optimizers empirically outperform their first-order counterparts in both speed and accuracy, the convergence theory for training such models is essentially confined to gradient descent. In this paper, we close this gap and develop a convergence analysis for the regularized Newton as well as Gauss-Newton method. As the number of hidden units tends to infinity, the training dynamics converge in probability, with explicit rates, to the solution of a deterministic limit equation involving the Newton neural tangent kernel associated with the PDE operator. Although the limiting kernel admits no spectral gap, we prove that, in the overparameterized limit, the PDE residual vanishes as training proceeds, and, under mild assumptions, the neural network converges to the PDE solution. We further show that, provided a sufficiently small regularization parameter, the projection of the residual onto any fixed number of leading eigenfunctions decays exponentially fast, demonstrating that this hyperparameter acts as a spectral filter: it lifts the convergence rates of high-frequency components towards those of low-frequency ones, thereby mitigating the spectral bias inherent in gradient descent. Mathematical challenges arise from the implicitly defined parameter update, involving a potentially indefinite Hessian, and from the PDE operator entering the kernel.

Part of the parallel session: “Neural Network approximations and related techniques for high-dimensional PDEs and Control Problems”.

Session 15D: Optimisation challenges in modern mobility *Room: Lise Meitner (A1.116)*

Large-Scale Analysis of Multi-Scale Queuing Networks for Car-Sharing Systems

Alessia Rigonat, Christine Fricker

INRIA Paris, DI ENS

Keywords: Stochastic networks, Fluid limits, Car sharing, Random environment

We develop a probabilistic model for free-floating car-sharing systems. Our goal is to characterize the macroscopic behavior of the system and provide insights for optimizing vehicle distribution. In these systems, shared vehicles occupy the same parking spaces as private cars. The availability of parking spaces across different zones of a city depends on private cars, which are far more numerous than free-floating vehicles. The system dynamics are described by a closed network of queues in which private and free-floating cars move among the same nodes, representing city zones. Nodes have finite capacity, and saturation is handled through a blocking and rerouting policy. The model accounts for spatial heterogeneity in both user demand and availability of parking spaces. We show that these dynamics preserve the product-form structure of the invariant distribution. We identify, in this setting, phase transitions between an overloaded regime where all nodes are saturated and underloaded regimes with saturated and non-saturated nodes. Scaling limits and stochastic averaging methods are used to analyze the behavior of the system when the capacity of the nodes is large. The analysis is performed when the average number of private cars per zone increases linearly with capacity, while the number of free-floating cars remains of smaller order.

Session 8J: Optimal Tariffs *Room: Emmy Noether (A1.140)*

Competitive Menu Pricing for Electric Vehicle Charging Stations

Rita Safi, Margarida Carvalho, Yezekael Hayel, Raphaël Payen, Julien Pennec, Tania Jiménez

1EDF R&D, France, rita.safi@edf.fr, raphael.payen@edf.fr, julien.pennec@edf.fr 2Université de Montréal, Canada, carvalho@iro.umontreal.ca 3Laboratoire Informatique d'Avignon, France, yezekael.hayel@univ-avignon.fr, tania.jimenez@univ-avignon.fr

Keywords: Electric vehicle, charging station, pricing competition, best-response algorithm

This work investigates a pricing competition framework among multiple Charging Stations (CSs). Each CS acts as a strategic player aiming to maximize its profit by attracting Electric Vehicles (EVs) through personalized charging offers. EV users select a CS based on both the proposed charging price and the travel distance to the station in order to maximize their utility. CSs seek flexibility from EV users by encouraging delayed charging schedules, enabling them to shift charging to periods of lower electricity prices. Accordingly, each EV is offered multiple charging options, such as immediate or flexible charging, each associated with different prices. Therefore, EV users must decide not only which charging station to visit but also which charging option to choose. To resolve situations in which multiple charging stations provide the same utility to an EV, a ranking preference mechanism is introduced. The resulting interaction is modeled as a game-theoretic framework, and the uniqueness of the equilibrium as well as the convergence of the Best-Response Dynamics algorithm are analyzed.

Session 15A: Optimal Control, Measures and PDE constrained optimization *Room: Marie Curie (Auditorium)*

Wasserstein-Fisher-Rao differential calculus and optimal control of weighted particle systems

Enrico Sartor[1]

[1]Laboratoire des Signaux et Systèmes, CNRS, CentraleSupélec, Université Paris-Saclay,
enrico.sartor@centralesupelec.fr

Keywords: Wasserstein-Fisher-Rao geometry, differential calculus on measure spaces, unbalanced optimal transport, weighted particle systems, optimal control

In this work, we introduce a static notion of differentiability for functionals defined on the space of positive measures, based on the Wasserstein-Fisher-Rao geometry, which accounts simultaneously for mass transport and mass creation or destruction [1, 2]. We establish general differentiability criteria and derive explicit formulas for relevant classes of functionals, including expectations, cylindrical functions, and interaction energies. A key feature of the resulting calculus is that differentiability along optimal Wasserstein-Fisher-Rao perturbations yields robust first-order expansions along more general, non-optimal perturbations. On the space of probability measures, our framework naturally connects with both Wasserstein differentiability and flat differentiability, providing a common geometric perspective on these two classical notions. Finally, we show how this differential calculus can be used in the optimal control of interacting particle systems with time-dependent weights, where both particle positions and weights are dynamical variables.

References

[1] Léna¨c Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. An interpolating distance between optimal transport and Fisher-Rao metrics. Foundations of Computational Mathematics, 18(1):1–44, 2018.

[2] Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new Hellinger-Kantorovich distance between positive measures. Inventiones mathematicae, 211(3):969–1117, 2018.

Session 5H: The moment hierachy and SDP *Room: Alan Turing (A1.134)*

Spectral-Gauge Cuts for Semidefinite Optimization

Antonio Sasaki[1], Sophie Demassey[1], Valentina Sessa[1]

[1]Centre de Mathématiques Appliquées (CMA), Mines Paris – PSL, France,
{antonio.sasaki,sophie.demassey,valentina.sessa}@minesparis.psl.eu

Keywords: semidefinite programming, cutting-plane methods, symmetric gauges, boxconstrained quadratic programming, sparse principal component analysis

Semidefinite relaxations are among the strongest tools available for some difficult problems, but become costly as the dimension grows. Cutting-plane methods address this bottleneck by starting from a cheaper outer approximation and enforcing positive semidefiniteness progressively through valid inequalities. We place this strategy in a general framework by normalizing positive semidefinite separators with a symmetric gauge, i.e., norms invariant under permutations and sign changes of the coordinates. Given a symmetric matrix $\bar{X}$ and a gauge ϕ, the resulting separation problem admits the closed-form optimal value

$$\min_{S \succeq 0,\ \phi(\lambda(S)) \leq 1} \langle S, \bar{X} \rangle = -\phi^{\circ}(\lambda_{-}(\bar{X})),$$

where ϕ° is the dual gauge and $\lambda_{-}(\bar{X})$ is the negative spectral part. This approach extends and unifies cut constructions that had previously appeared separately. The $\ell 1$ gauge recovers the classical eigencut introduced by Ramana [1] and later used with the reformulation-linearization technique by Sherali and Fraticelli [2]. Support restrictions give the sparse eigencuts of Dey et al. [3], while the $\ell\infty$gauge yields the nuclear cuts of Bertsimas and Cory-Wright [4]. Beyond these special cases, within the ℓp norms, each $1 < p < \infty$defines a new intermediate cut. More generally, gauges can be designed to enforce properties such as low rank or restricted support, yielding families with different tightness-cost trade-offs. We embed the oracle in a Kelley-type algorithm, prove convergence to the semidefinite optimum, and test it on instances of box-constrained quadratic programming (BoxQP) and sparse principal component analysis (sparse PCA).

References

[1] M. V. Ramana, An algorithmic analysis of multiquadratic and semidefinite programming problems, Ph.D. thesis, Johns Hopkins University, 1993.

[2] H. D. Sherali and B. M. P. Fraticelli, Enhancing RLT relaxations via a new class of semidefinite cuts, J. Global Optim., 22:233–261, 2002.

[3] S. S. Dey, A. M. Kazachkov, A. Lodi, and G. Muñoz, Cutting plane generation through sparse principal component analysis, SIAM J. Optim., 32(2):1319–1343, 2022.

[4] D. Bertsimas and R. Cory-Wright, On polyhedral and second-order cone decompositions of semidefinite optimization problems, Oper. Res. Lett., 48(1):78–85, 2020.

Session 15E: Algorithms 2.0 : advances in computational methods *Room: Antoine Lavoisier (A1.122)*

Solving Minimum-Cost Flows on GPUs

Simeon Schaub, Guillaume Dalle

1CERMICS, CNRS, ENPC, Institut Polytechnique de Paris, Marne-la-Vallée, France, simeon.schaub@enpc.fr, 2LVMT, Univ Gustave Eiffel, ENPC, Institut Polytechnique de Paris, Marne-la-Vallée, France guillaume.dalle@enpc.fr

Keywords: Network Flows, Primal-Dual Hybrid Gradient, Sparse Linear Algebra, GPUs

Primal-dual hybrid gradient (PDHG) and its variants [1] have proven to be efficient for solving large-scale linear programs, thanks to the parallelism enabled by GPUs. Since integer flows can be modeled as linear programs, this begs the question whether primal-dual methods are suitable for minimum cost flow problems. We assume multiple similar instances need to be solved at the same time, as is common when training differentiable optimization layers, and investigate the competitiveness of PDHG against specialized flow solvers. Prior work focused on using PDHG for solving large multicommodity flow problems [2], but to the best of our knowledge, general minimum cost flow problems have not been tackled before. We will use the problems posed in the 13th DIMACS challenge on network flows as our benchmarks1. Since sparse matrix-vector products represent the main bottlenecks of PDHG, we aim to accelerate these products by taking advantage of two special properties of the node-arc incidence matrix: there are exactly two entries per column, and arcs can always be sorted according to their origin node, putting half of the coefficients in a straightforward staircase structure. Handling multiple similar instances at the same time allows us to use batching [3]: Sparse matrix-vector products become matrix-matrix products, which improves throughput. Our preliminary results suggest that taking advantage of special properties of the incidence matrix alone does not bring significant benefits, but they reveal other ways of speeding up existing PDHG implementations.

References

[1] David Applegate et al. “Practical large-scale linear programming using primal-dual hybrid gradient”. In: Proceedings of the 35th International Conference on Neural Information Processing Systems. Curran Associates Inc., 2021.

[2] Fangzhao Zhang and Stephen Boyd. Solving Large Multicommodity Network Flow Problems on GPUs. Apr. 3, 2025. arXiv: 2501.17996[math.OC].

[3] Nicolas Blin et al. “Batched First-Order Methods for Parallel LP Solving in MIP”. In: Forty-third International Conference on Machine Learning. June 24, 2026.

1https://coral.ise.lehigh.edu/flow-challenge-2-0/mcf-instances-and-generators/

Session 12H: Market Design and Strategic Behavior *Room: Alan Turing (A1.134)*

SS-TEA: Search-Supervised Transformer Edge Addition for Black-Box Graph Intervention

Xu Shuangshuang (Mme)[1], Andrea Araldo[2], Duo WANG[3]

[1]École Polytechnique, Institut Polytechnique de Paris, France, shuangshuang.xu@polytechnique.edu
[2]Télécom SudParis, Institut Polytechnique de Paris, France, andrea.araldo@telecom-sudparis.eu
[3]Peking University, China, 2606397279@pku.edu.cn

Keywords: graph intervention, black-box optimization, Transformer, combinatorial optimization, cross-size generalization

Across transportation, communication, and graph-based physical systems, topology determines how traffic and information flow and how physical components interact, thereby shaping systemlevel behavior and robustness. Consequently, adding only a small number of edges can substantially alter graph-level outcomes. However, identifying effective edge additions is challenging: the candidate space grows quadratically with graph size, while the objective is often accessible only through an expensive and non-differentiable black-box solver or simulator. Existing learningbased approaches mainly formulate edge selection as sequential reinforcement learning, often requiring long training and scale-specific optimization. In this work, we propose Search-Supervised Transformer Edge Addition (SS-TEA), a unified Transformer framework for budgeted graph edge addition under black-box feedback. SS-TEA captures global node dependencies through full self-attention and uses a masked query–key decoder to score feasible nonedges, constructing a complete fixed-budget edge set from a single model forward pass. To train the Transformer without evaluator gradients, high-quality edge sets discovered during search are converted into supervision for the edge-selection policy. Extensive experiments demonstrate that SS-TEA achieves strong performance with a small training budget and transfers directly from small to large graphs without retraining. The same framework applies to both exposing solver vulnerabilities and improving network robustness, demonstrating its cross-scale transferability and broad cross-task applicability.

Session 15I: Geometry *Room: Sophie Germain (A1.139)*

Exact Extremal Optimization for Wrong-Sign Entries in Equal-Diagonal Symmetric Matrices

Harman Singh

University of Waterloo, Canada

Keywords: matrix optimization, spectral optimization, nonconvex optimization, symmetric matrices, correlation matrices

We study an exact global extremal problem on nonconvex equal-diagonal isospectral families of symmetric matrices. Let $A \in \mathbb{R}^{n\times n}$ be symmetric with constant diagonal d, simple largest eigenvalue λ, and remaining eigenvalues in $[a, b]$, where $b < \lambda$. If v is a top eigenvector, an off-diagonal entry A_{ij} is wrong-sign when $A_{ij}v_iv_j < 0$. For a prescribed mean-d spectrum Σ with simple maximum λ, second-largest member b, and smallest member a, we maximize $|A_{12}|$ over the equal-diagonal isospectral fiber subject to $v_\lambda(A) > 0$ and $A_{12} < 0$. We determine the supremum exactly. If $b \le d$, no wrong-sign pair is feasible. If $b > d$, the optimal value is

$$\tau_d(\lambda; a, b) = \begin{cases} \sqrt{\dfrac{(d-a)(b-d)(\lambda-d)}{\lambda-a-b+d}}, & a+b < 2d, \\ d-a, & a+b \ge 2d, \end{cases}$$

and the supremum is not attained under the strict positivity constraint. In the first regime, this value is strictly below both classical interlacing caps, so the sign constraint yields a genuine improvement over sign-blind spectral bounds. The upper bound follows from complementary positive-semidefinite spectral residuals, while global sharpness is obtained using Schur–Horn realizations and diagonal-preserving deformations inside the isospectral fiber. Fixing the dimension and optimizing the remaining spectral freedom gives a second exact extremal value and an exact feasibility condition. For correlation matrices, the resulting constant is explicit in λ_2, λ_N, and the dimension.

Session 12J: Handling dynamics *Room: Emmy Noether (A1.140)*

Monotone Inclusions for Decision-Dependent Stochastic Problems

Wutao Si, Hamza Ennaji, Jalal Fadili

1 Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France wutao.si@univ-littoral.fr, hamza.ennaji@univ-grenoble-alpes.fr 2 ENSICAEN, Normandie Université, CNRS, GREYC, France jalal.fadili@ensicaen.fr

Keywords: Monotone inclusions, decision-dependent distributions, performative prediction, forward-backward splitting, inertial dynamics

Many learning and decision-making problems are solved in an environment that reacts to the decision itself: the data distribution is not fixed, but is induced by the deployed decision. This is the setting of stochastic optimization with decision-dependent distributions, also known as performative prediction, where a natural solution concept is an equilibrium, i.e. a decision which is optimal for the distribution it generates [2]. In a real Hilbert space H, this leads to the closed-loop monotone inclusion

$$\text{find } \bar{x} \in \mathcal{H} : \qquad 0 \in A(\bar{x}) + B_{m_{\bar{x}}}(\bar{x}).$$

where A is maximally monotone, Bm(x) = Eξ∼m[B(x, ξ)] is cocoercive, and (mx)x∈H is a family of probability measures indexed by the decision x. Taking A = ∂g and B(x, ξ) = ∇xf(x, ξ) recovers regularized stochastic optimization. This framework was introduced in [1], where firstand second-order dynamics are analysed through a reformulation as a Lipschitz perturbation of a monotone inclusion. Building upon this framework, we introduce second-order splitting dynamics [3] driven by a distributionally evaluated forward-backward operator. This approach is particularly efficient as it requires only forward evaluations of Bm and resolvents of A. Under a uniform monotonicity assumption, the equilibrium exists and is unique. Furthermore, a vanishing viscous damping yields strong convergence of the trajectories along with fast rates for the velocities, whereas a constant Polyak-type damping ensures global exponential convergence when the regularizing operator is strongly monotone.

References

[1] H. Ennaji, J. Fadili, H. Attouch, Stochastic monotone inclusion with closed loop distributions, Evolution Equations and Control Theory, 17 (2026), 140–172.

[2] D. Drusvyatskiy, L. Xiao, Stochastic optimization with decision-dependent distributions, Mathematics of Operations Research, 48(2):954–998, 2023.

[3] W. Si, H. Ennaji, J. Fadili, Second order splitting dynamics for stochastic monotone inclusions with closed loop distribution, 2026, arXiv:2606.06280.

Session 5D: Invited session: Black-box Optimization *Room: Lise Meitner (A1.116)*

Blackbox Optimization for Active Distribution Networks Reconfiguration under Uncertainty

Christina G. Soldati, Hussein Suprême, Miguel F. Anjos, Sébastien Le Digabel, Antoine Lesage-Landry

1Polytechnique Montréal, Canada, christina-g.soldati@etud.polymtl.ca, sebastien.le-digabel@polymtl.ca, antoine.lesage-landry@polymtl.ca 2Groupe d'études et de recherche en analyse de décisions (GERAD), Canada 3School of Mathematics, Maxwell Institute for Mathematical Sciences, University of Edinburgh, UK anjos@stanfordalumni.org 4MILA (Institut québécois d'intelligence artificielle – Quebec Artificial Intelligence Institute), Canada 5 Institut de recherche d'Hydro-Québec (IREQ), Canada, supreme.hussein2@hydroquebec.com

Keywords: Blackbox Optimization, Distributed Energy Resources, Power Distribution Network, Reconfiguration, Stochastic Optimization

The energy transition is reflected in part by the modernization of distribution networks (DNs), which are evolving toward a more active role in power system operations. These active distribution networks (ADNs) are characterized by a high penetration of distributed energy resources (DERs), often of renewable sources, and a flexible network topology enabled by remotely controlled switches. DERs can induce bidirectional power flows that may exacerbate power losses, cause voltage violations, lead to line overloading, and affect frequency stability, among others. DNR can be modelled as a mixed-integer nonlinear optimization problem that seeks a radial topology to minimize power losses through switch operations. The proposed approach for DER integration using DNR combines blackbox optimization (BBO) and a high-accuracy power flow solver for improved feasible topology identification. Its computational burden is reduced by splitting the problem into continuous and binary formulations, solved respectively by BBO and a heuristic combinatorial optimization algorithm adapted to the DN context. To account for uncertainty in load demand and DER generation, the approach is formulated as a two-stage scenario-based stochastic program minimizing expected day-ahead power losses. In the first stage, the heuristic explores the binary decision space to identify a topology. Given this topology, the second stage solves a continuous BBO sub-problem for each scenario, where DER generations provide the recourse to ensure the topology feasibility across all considered scenarios. Each scenario combines load demand and DER generation limits from the same hour, preserving temporal correlation, and is derived from historical demand profiles and solar irradiance data. Spatial correlation is preserved through scaling according to the test DN baseline values. To reduce the computational burden of the second stage, scenario reduction is performed to identify representative clusters of similar scenarios. Inspired by multi-fidelity optimization, the fidelity level is defined by the number of clusters considered and is gradually increased at each iteration. One scenario is randomly sampled from each cluster. Combined with a reduced BBO evaluation budget, it mitigates the computational effort in the early iterations. The proposed method is then evaluated on three DNs of increasing size to assess its scalability.

Session 10C: Combinatorial optimization augmented machine learning *Room: Henri Poincaré (Amphi I)*

Combinatorial Optimization Augmented Machine Learning for Dynamic Electric Autonomous Dial-a-Ride Problem

Yue Su[1], Léa Ricard[2]

[1]Centre INRIA de l'Université de Lille, France, first.author@wonderland.fr
[2]École Polytechnique Fédérale de Lausanne, Switzerland, lea.ricard@epfl.ch

Keywords: Dial-a-ride problem, Electric autonomous vehicles, Decision-focused learning

This study introduces a decision-epoch-based dynamic electric autonomous dial-a-ride problem (Dyn-EADARP), in which incoming requests are collected and processed at periodic decision epochs. A key decision is not only how to serve requests, but also when to dispatch them. At each decision epoch, the service provider decides which requests to serve and which to postpone, while jointly determining vehicle routes, schedules, and charging decisions. Unexecuted parts of existing plans can be revised as new information becomes available. To solve this problem, we develop an ML–CO policy following the combinatorial optimization augmented machine learning (COAML) framework, which combines a statistical model with a combinatorial optimization layer for decision making. The statistical model predicts prizes for available requests, and a prize-collecting E-ADARP uses these prizes to jointly determine request selection, routing, scheduling, and charging. The statistical model is trained directly to improve the decisions produced by the optimization layer. Computational experiments on 320 test instances demonstrate the efficiency of ML–CO, which solves instances with nearly 500 requests in about 6 seconds on average. It achieves 9.8%–16.8% lower objective values than benchmark policies and an average gap of 4.8% to the anticipative reference. The results provide several managerial insights. First, serving requests immediately is not always best, as selectively postponing some requests can create better ride-sharing opportunities. Second, revising existing plans preserves operational flexibility and substantially improves solution quality. Finally, more frequent decision making does not necessarily improve performance, highlighting the importance of choosing an appropriate decision frequency.

Session 12E: Convergence rates *Room: Antoine Lavoisier (A1.122)*

Convergence rates for the moment-SoS hierarchy

Corbinian Schlosser[1], Matteo Tacchi-Bénard[2], Alexey Lazarev[3]

[1]INRIA Paris, France, corbinian.schlosser@gmail.com
[2]CNRS, GIPSA-lab, France, matteo.tacchi@gipsa-lab.fr,
[3]Institut de Mathématiques de Toulouse, France, m_laz1@mail.ru

Keywords: Nonconvex optimization, SoS programming, convex embeddings, Borel measures, real algebraic geometry

Since the early 2000s, real algebraic geometry found a new application in the moment- SoS hierarchy, a framework that leverages Putinar's Positivstellensatz to solve generalized moment problems (GMP). This methodology was first applied to polynomial optimization, before finding new applications in optimal control, volume computation, option pricing, stability analysis, study of stochastic processes, solving PDEs and super resolution, to name a few. The moment-SoS hierarchy relies on positivity certificates, that consist in polynomial sums of squares of bounded degree. Finding the maximum degree bound required to certify positivity of a given polynomial has been an important research topic, to which a recent contribution, namely the Effective Putinar Positivstellensatz [1], provided a promising answer. Based on this, it is possible to compute the convergence rate of the moment-SoS hierarchy, depending on the considered GMP instance [2, 3]. In this talk, I will explain this new Effective Putinar Positivstellensatz, and illustrate its application to computing the convergence rate of some key applications of the moment-SoS hierarchy.

References

[1] L. Baldi and B. Mourrain, On the effective Putinar's Positivstellensatz and moment approximation, Mathematical Programming A, 200:71–103, 2022.

[2] C. Schlosser, M. Tacchi and A. Lazarev, Convergence rates for the moment-SoS hierarchy, Numerical Algebra, Control and Optimization, 16:105–156, 2026.

[3] C. Schlosser and M. Tacchi, Specialized Positivstellensätze for improved convergence rates of the moment-SoS hierarchy, IEEE Control Systems Letters, 8:2319–2324, 2024.

Session 12C: Invited session: Geometric Control and Optimization *Room: André Marie Ampère (Amphi II)*

Branching of Regular Abnormal Geodesics in Sub-Riemannian Manifolds of Rank Two

Luca Rizzi[1], Mario Sigalotti[2], Nicolò Tedesco[1,2]

[1]SISSA, Trieste, Italy, lrizzi@sissa.it, ntedesco@sissa.it
[2]Sorbonne Université, Inria, Laboratoire Jacques Louis-Lions, Paris, France, mario.sigalotti@inria.fr
Invited session: Geometric Control and Optimization

Keywords: Control-linear systems, Branching of optimal trajectories, Abnormal minimizers.

When considering an optimal control problem, we say that two distinct extremals branch if, up to time translations, they coincide on a subinterval. Our goal is to describe branching properties of optimal trajectories for a class of optimal control problems with cost $J(u) = \int |u|^2\, dt$, where u denotes the control. We consider nonlinear control systems in $\mathbb{R}^n$, or more generally on a manifold M, with dynamics that are linear in the controls. By introducing the endpoint map it is possible to rephrase the optimal control problem as a constrained minimization problem on L^2 and deduce, as an application of the Lagrange multipliers rule, the existence of two classes of candidate optimal trajectories: normal and abnormal extremals. Normal extremals are smooth and can be lifted to solutions of a Hamiltonian system on T^*M. Abnormal extremals do not depend on the particular choice of the cost and may be nonsmooth. However, their presence must be taken into account since there exist optimal trajectories that are strictly abnormal (i.e. they are not also normal).

In [1], it was proven that normal extremals can branch, and a characterization of the branching times was provided. Our study involves a relevant class of smooth abnormal extremals for systems with two-dimensional controls, called regular [2]. These extremals, analogously to normal ones, can be lifted to solutions of a Hamiltonian system on a submanifold $R \subset T^*M$. This allows to adapt the arguments of [1], obtaining a characterization of their branching times which is formally analogous to the one given in [1] for normal extremals. We also discuss two cases in which the branching curves are strictly abnormal; in one of them, we consider a reference path that is both normal and regular abnormal, that splits at the same time into two distinct families of branching curves that are, respectively, normal and regular abnormal, proving that the regular abnormal ones are also strictly abnormal.

Finally, we will discuss examples of extremals with an arbitrarily large number of branchings, possibly accumulating at finite time.

References

[1] T. Mietton and L. Rizzi. Branching geodesics in sub-Riemannian geometry. Geom. Funct. Anal., 30(4):1139–1151, 2020

[2] W. Liu and H. J. Sussman. Shortest paths for sub-Riemannian metrics on rank-two distributions. Mem. Amer. Math. Soc., 118(564):x+104, 1995

Session 8B: Invited session: Linearly Controlled Systems *Room: Henri Poincaré (Amphi I)*

Metric limits of sub-Finsler Lie groups

Enrico Le Donne[1], Sebastiano Nicolussi Golo[2], Andrea Tettamanti[1]

[1]University of Fribourg, Switzerland, enrico.ledonne@unifr.ch, andrea.tettamanti@unifr.ch
[2]sebastiano2.72@gmail.com

Keywords: Sub-Finsler geometry, Lie groups, Asymptotic geometry

We study the small and large scale geometry of sub-Finsler metrics on Lie groups. A subFinsler distance can be viewed as the value function of a minimum-time optimal control problem for a driftless system that is linear in the controls. By two theorems of Mitchell and Pansu, under suitable hypotheses, both the tangent cone and the asymptotic cone of Lie groups equipped with left-invariant sub-Finsler metrics are Carnot groups. In this talk, we present quantitative estimates on the rate of convergence to these limiting spaces, obtained in [1].

References

[1] E. Le Donne, S. Nicolussi Golo, A. Tettamanti, Quantitative versions of Pansu Asymptotic Theorem and of Mitchell Tangent Theorem, arXiv:2601.21509.

Session 5F: IS : Recent Advances in Systematic and Computer-Aided Methods in Optimization I *Room: Ada Lovelace (A1.128)*

Computer-Aided Proofs in First-Order Optimization, with Applications to Error Feedback

D. B. Thomsen, A. Taylor, A. Dieuleveut

Keywords: First-order optimization, performance estimation, computer-aided proofs, error feedback, distributed optimization

First-order methods are widely used in optimization and machine learning, and their behavior is often analyzed through the spectrum of worst-case convergence rates. Obtaining such guarantees is often difficult, time-consuming, and error-prone. Starting with the work of Drori and Teboulle (2014), novel techniques have been used to gain numerical insights, leading to the release of various performance estimation (PE) software. In this talk, I will show how various computer-aided techniques can be used to study firstorder optimization methods in a systematic way. From performance estimation problems with automated Lyapunov discovery to symbolic regression and computer algebra systems, novel tools completely reshape the way we approach the theory of optimization. As a main example, I will focus on error feedback methods used with compressed communication in distributed optimization. While error feedback has been widely studied, existing theory often provides non-tight, and thus unreliable, bounds. I will present tight analyses with matching lower bounds [2, 1] that allow a fair comparison between error feedback schemes and standard compressed gradient descent, and help explain when error feedback is useful and when it is not. Overall, the talk aims to show how various computer-aided proofs can lead to clearer and more reliable insights into first-order optimization methods.

References

[1] D. B. Thomsen, A. Taylor, and A. Dieuleveut, A Tight Theory of Error Feedback Algorithms in Distributed Optimization, International Conference on Machine Learning (ICML), 2026. Conference page.

[2] D. B. Thomsen, A. Taylor, and A. Dieuleveut, Tight Analyses of First-Order Methods with Error Feedback, Advances in Neural Information Processing Systems (NeurIPS), 2025.

[3] B. Goujaud, C. Moucer, F. Glineur, J.-M. Hendrickx, A. B. Taylor, and A. Dieuleveut, PEPit: Computer-Assisted Worst-Case Analyses of First-Order Optimization Methods in Python, Mathematical Programming Computation, 16(3):337–367, 2024.

[4] B. Goujaud, A. Taylor, and A. Dieuleveut, Provable Non-Accelerations of the Heavy-Ball Method, Mathematical Programming, 2025.

Session 8D: Invited session: Recent Advances in Systematic and Computer-Aided Methods in Optimization part 2
Room: Lise Meitner (A1.116)

Simple Proofs, and How to Find Them

Daniel Berg Thomsen, Manu Upadhyaya, Baptiste Goujaud, Aymeric Dieuleveut, Adrien Taylor

[1]Inria & D.I. École Normale Supérieure, CNRS & PSL Research University, Paris, France; [2]CMAP, École Polytechnique, Institut Polytechnique de Paris, France; [3]SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, France

Keywords: First-order optimization, performance estimation, proof complexity, simple proofs

Performance estimation has put first-order optimization in a remarkable position: for many algorithms and problem classes, we can now compute tight worst-case convergence rates and extract numerical certificates of optimality. The harder question is no longer whether a proof exists, but whether it can be made simple. In many cases, the available certificate is a sum-of- squares identity with coefficients so intricate that checking the algebra by hand is itself a test of patience. This talk introduces several notions of proof complexity for first-order methods and asks how they can be used to search for simpler proofs. The main idea is to exploit the proof structure of problems in first-order optimization, together with heuristics inspired by sparse optimization, to turn numerical certificates into short, interpretable arguments. We also discuss methods that automatically recover the simplest proof within a chosen proof class, along with examples. The goal is to bridge the gap between computational proof discovery and human- readable mathematics.

Session 15G: Data Science, Learning and the intersection with Optimization *Room: Enrico Fermi (A1.133)*

Sensitivity of GPU Scheduling Rules to Runtime Estimation Error

Yosef Gavriel Tirat-Gefen, Filipe Vinci

1 MaxWave Research LLC, Rockville, Maryland, USA; yosefgavriel@gmail.com 2 CentraleSupélec - Campus Paris-Saclay, Université Paris-Saclay, France

Keywords: GPU scheduling; uncertain processing times; deadlines; simulation

GPU schedulers must allocate limited capacity among jobs with different deadlines and uncertain runtimes. Earlier work has studied deadline-aware GPU scheduling [1] and robust use of runtime predictions in cluster scheduling [2]. We examine how estimation error changes the performance of simple priority rules.

We simulate 80 jobs competing for four identical GPUs. Four non-preemptive rules use the same backfill procedure: first-come-first-served, earliest-deadline-first (EDF), estimated slack, and buffered slack. For a waiting job j at time t, estimated slack is $d_j - t - r_j$, where d_j is its deadline and r_j its estimated runtime. Buffered slack instead uses

$$s_j(t) = d_j - t - r_j \exp(1.28\sigma_j),$$

where σ_j specifies the simulated level of runtime-estimate error. Jobs with the smallest slack receive priority. We compare the rules on paired synthetic job streams while varying arrival rate, deadline tightness, and multiplicative estimation error.

Across 100 trials per setting, no rule dominates. Under heavy load with exact estimates, buffered slack missed 54.80 deadlines per trial versus 51.96 for EDF. Under medium load, tight deadlines, and high estimate noise, buffered slack missed 38.54 versus 40.79 for EDF. A second set of 100 trials with a different random seed reproduced the latter reversal.

These results illustrate how runtime-estimate uncertainty can change the performance of a priority rule. The findings are limited to synthetic jobs and a specified error model; evaluation on measured GPU workloads remains necessary.

References

[1] W. Gao et al., “Chronus: A Novel Deadline-aware Scheduler for Deep Learning Training Jobs,” ACM SoCC, 2021. doi:10.1145/3472883.3486978.

[2] A. Tumanov et al., “JamaisVu: Robust Scheduling with Auto-Estimated Job Runtimes,” Carnegie Mellon University, 2016. Technical report.

Session 5I: Invited Session:Neural Network approximations for high-dimensional PDEs Problems*Room: Sophie Germain (A1.139)*

When does residual minimization solve a differential equation?

Richard Tsai, Carlos Esteve-Yagüe, Stanley Osher

University of Texas at Austin; University of Alicante; UCLA

Keywords: least squares, physics-informed neural networks, residual minimization, parameterized trial spaces

Least-squares residual minimization approximates a differential equation through an optimization problem. This paper examines the implications needed to pass from stationary points and minimizing sequences of the residual functional to the intended solution. Starting from classical range–adjoint geometry, we study how boundary conditions and parameterized variations determine what can be inferred about the residual at a stationary point. Smooth spurious critical points can occur even for well-posed problems; for linear equations, a compatibility condition yields an exact projection characterization. For piecewise-smooth trial functions, we identify conditions under which the mean gradient of the sampled loss differs from the continuous-loss gradient, including an explicit example with identically vanishing sampled gradient and nonzero continuous derivative. We then examine two further issues: for hyperbolic conservation laws, transonic minimizing sequences can lose regularity and force parameter escape, while an exactly vanishing Hamilton–Jacobi residual does not necessarily lead to the viscosity solution. Favorable results for a viscous model and uniformly convex Monge–Ampère critical points identify conditions that restore the corresponding implications. Together, these results distinguish the roles of stationarity, sampling, compactness, and solution admissibility in interpreting residual minimization, with particular relevance to physics-informed neural networks.

Session 8B: Invited session: Linearly Controlled Systems *Room: Henri Poincaré (Amphi I)*

Pontryagin Maximum Principle for Non-Smooth State-Constrained Control Problems over Wasserstein Spaces

Fernanda Urrea[1], Ernesto Treumún[2]

[1]INSA Rouen Normandie, France, fernanda.urrea@insa-rouen.fr
[2]ENSTA - Institut Polytechnique de Paris, France, ernesto.treumun@ensta.fr

Keywords: Pontryagin maximum principle, Clarke subdifferential, Continuity equation, Wasserstein spaces, State constraints

In this talk, we study Pontryagin-type optimality conditions for optimal control problems of non-local continuity equations over the Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$. Whenever the underlying data is smooth, the theory for addressing these problems is well-established in the literature. However, the class of smooth maps over $\mathcal{P}_2(\mathbb{R}^d)$ is rather scarce. In fact, even the squared Wasserstein distance fails to be smooth unless strong assumptions are imposed on the measures.

We overcome this difficulty by introducing a new map-based subdifferential on $\mathcal{P}_2(\mathbb{R}^d)$, in the spirit of Clarke's subdifferential in Banach spaces. Using this new object, we establish a Pontryagin maximum principle for state-constrained control problems in which both the cost and constraint functionals are only locally Lipschitz.

References

[1] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. Gradient Flows: In Metric Spaces and in the Space of Probability Measures. Lectures in Mathematics. ETH Zürich Ser. Springer Basel AG, Basel, 1st edition, 2005.

[2] Benoît Bonnet and Hélène Frankowska. Necessary optimality conditions for optimal control problems in wasserstein spaces. Applied Mathematics and Optimization, 84, 2021.

[3] Fernanda Urrea and Ernesto Treumún. Pontryagin maximum principle for non-smooth stateconstrained control problems over wasserstein spaces, 2026. arXiv:2609.17197.

Session 15A: Optimal Control, Measures and PDE constrained optimization *Room: Marie Curie (Auditorium)*

Sensitivity Analysis for free terminal time horizon problems in Wasserstein spaces

Fernanda Urrea[1]

[1]INSA Rouen Normandie, France, fernanda.urrea@insa-rouen.fr

Keywords: Wasserstein spaces, Optimal control, Sensitivity analysis, Pontryagin maximum principle, Free-time horizon, Minimal time problems.

In this talk, we study free-terminal-time optimal control problems for non-local continuity equations over the Wasserstein space ($\mathcal{P}_2(\mathbb{R}^d)$). Although Pontryagin's maximum principle and dynamic programming provide complementary descriptions of optimal trajectories, connecting them through sensitivity relations is delicate in this setting. In particular, the value function is generally nonsmooth, and subgradients computed at different probability measures belong to different L^2 spaces. Under a Petrov-type condition on the terminal target, we obtain a quantitative steering estimate and prove local Lipschitz continuity of the value function around an optimal trajectory. We then adapt the classical auxiliary-problem approach to the Wasserstein framework and establish the existence of a Pontryagin coextremal satisfying both full time–measure and partial measure sensitivity relations.

References

[1] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. Gradient Flows: In Metric Spaces and in the Space of Probability Measures. Lectures in Mathematics. ETH Zürich Ser. Springer Basel AG, Basel, 1st edition, 2005.

[2] Frank H Clarke and Richard B Vinter. The relationship between the maximum principle and dynamic programming. SIAM Journal on Control and Optimization, 25(5):1291–1311, 1987.

Session 8H: Learning dynamics and stochasticity in optimization *Room: Alan Turing (A1.134)*

Implicit Bias of SGD for Overparametrized Models

Govind Menon[1], Austin J. Stromme[2], Adrien Vacher[2]

[1]Brown University, USA, govind.menon@brown.edu
[2]ENSAE, France, austin.stromme@ensae.fr, adrien.vacher@ensae.fr

Keywords: Implicit Bias, Langevin Dynamics, Isometric Group Actions, Geometry

Overparameterization is a defining feature of modern machine learning, but understanding how it affects the dynamics and implicit regularization of Stochastic Gradient Descent (SGD) remains a major challenge. In this talk, we study a tractable form of overparameterization arising from symmetries: different parameter values can represent the same function, and the corresponding transformations of the parameters can be described by an isometric group action. Such symmetries appear naturally in modern architectures, including transformers. We will first present recent work, published at COLT 2026 [1], in which we model SGD for functions with such symmetries by a continuous-time Langevin dynamics whose noise covariance is given by the orthogonal projection onto the normal space to the group orbit. Thus, the noise acts only in directions that modify the represented function, while directions tangent to the orbit corresponding to different parameterizations of the same function are not affected. We show that, when the initial and target distributions are invariant under the group action, this apparently anisotropic Langevin dynamics is equivalent in law to standard Langevin dynamics with isotropic noise, but with an additional drift term given by the gradient of the negative logarithm of the volume of the group orbit. In other words, the geometry of the parameter space itself induces an implicit regularization: the dynamics is biased towards regions where the symmetry orbits have smaller volume. We establish this result by coupling the dynamics with an auxiliary stochastic process evolving on the group, and identify the additional drift with the mean curvature of the group orbits. We will then turn to ongoing work on overparameterized matrix sensing, a classical toy model for non-convex learning. Here, a more refined description of the SGD noise reveals a strong numerical bias towards low-rank solutions. Interestingly, unlike the low-rank bias observed in gradient flow, this effect persists even when the initialization is not close to zero. Finally, we will discuss a broader question: how do different sources of overparameterization interact? In particular, can the geometric overparameterization induced by symmetries interact with the more familiar notion of having more parameters than training examples, and together explain some of the implicit regularization observed in SGD?

References

[1] G. Menon, A.J. Stromme, A. Vacher, On the implicit regularization of Langevin dynamics with projected noise, Conference On Learning Theory, 2026.

Session 12I: Optimization for Learning and Large-Scale Decision Making *Room: Sophie Germain (A1.139)*

Controlled Screening for Smooth Objectives

Eyal Vayness[1], Olivier Fercoq[2], Maxime Sangnier[1]

[1]Sorbonne Université, France, vayness@lpsm.paris, sangnier@lpsm.paris
[2]Télécom Paris, France, olivier.fercoq@telecom-paris.fr

Keywords: Sparsity, statistical learning, reproducing kernel Hilbert space

Screening is a technique introduced in [1], which leverages optimality conditions of a nonsmooth convex optimization problem to retrieve sparsity in features or samples. Most existing works are using non smooth objectives to guarantee null variables (via the loss for sample elimination [3] or on the penalty for feature sparsity [2] or both [4]). Safe screening based on duality gap bounds allows one to test during the training phase for relevance of optimization variables. In the context of reproducing kernel Hilbert spaces (RKHS), where the number of optimization variables is equal the number of samples, screening is a good way to relax the optimization problem on the go and reduce the problem's dimensions while being solved. However actual safe screening method that ensure nullity of optimization variables heavily rely on non smoothness. Given the same tools but with a smooth objective, we introduce here a way to screen variables that are not null but rather lie in a small interval near zero at the optimum. We place ourselves in the context of an RKHS H, a differentiable strictly convex loss and L^2 regularization. Finding the solution $h^\star \in H$ of the primal optimization is equivalent to finding the solution $\alpha^\star \in \mathbb{R}^n$ of the dual problem. A coordinate i for which $\alpha_i^\star = 0$ means that the i-th sample is irrelevant for the optimization problem and can be taken out. Applying standard safe screening tools, we provide sufficient conditions on any primal-dual candidate $(\widehat{h}, \widehat{\alpha})$, which ensures that $|\alpha_i^\star| < \delta$. After a coordinate has been screened one has two possibilities: (i) set $\widehat{\alpha}_i$ to zero which induces sparsity on the estimator and results in some δ-approximation of the real solution $\alpha^\star$ or (ii) leave the i-th coordinate as is and stop updating it with the optimization algorithm. Both solutions unburden the solver's next steps as one relaxed the problem to smaller dimensions.

References

[1] L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. Pacific Journal of Optimization, 8(4):667–698, 2012. [2] E. Ndiaye. Safe Optimization Algorithms for Variable Selection and Hyperparameter Tuning. Theses, Université Paris Saclay (COmUE), 2018. [3] K. Ogawa, Y. Suzuki, and I. Takeuchi. Safe Screening of Non-Support Vectors in Pathwise SVM Computation. In Proceedings of the 30th International Conference on Machine Learning, pages 1382–1390. PMLR, 2013. [4] A. Shibagaki, M. Karasuyama, K. Hatano, and I. Takeuchi. Simultaneous Safe Screening of Features and Samples in Doubly Sparse Modeling. In Proceedings of The 33rd International Conference on Machine Learning, pages 1577–1586. PMLR, 2016.

Session 8H: Learning dynamics and stochasticity in optimization *Room: Alan Turing (A1.134)*

Optimizing Factorable Losses via a Change of Variable: Properties and Computational Opportunities

Manon Verbockhaven, Arthur Lebeurrier, Rémi Gribonval

1ENS de Lyon, CNRS, Université Claude Bernard Lyon 1, Inria, LIP, author@ens-lyon.fr 2Inria, ENS de Lyon, CNRS, Université Claude Bernard Lyon 1, LIP author@inria.fr 3UMR 5668, 69342, Lyon cedex 07, France

Keywords: Optimization, Gradient Descent, Preconditioning, Approximation, Symmetries

Due to properties of the model parametrization or of the loss, many machine learning problems can be reformulated as the minimization of a composition of functions, $L = L \circ g$, where g is said to factorize the original objective function L. While such optimization problems are usually addressed in the domain of the original objective L, refereed to as the parameter space, we investigate the opportunities offered by a modification of the optimization domain using a change of variable with g, which also translates as the intent of optimizing the function L on the range of g rather than L on the parameter space. For an arbitrary g, we connect this approach with differential geometry through the Riemannian gradient of L on the manifold of the image of g, and identify the associated update of parameters as a preconditioned version of gradient descent, which, even for non-injective g as in the context of a quotient manifold, is associated with a well-defined Riemannian metric on the parameter space [1, 2]. While the computational challenges associated with such approaches are known in the literature namely, the computation of the preconditioner via the inverse of a large matrix S and its application to the usual gradient vector we show that the change of variable viewpoint offers a fresh perspective and opens new computational approaches. Indeed, it naturally allows direct approximation of S^{-1} via a criterion whose minimization can be performed without ever computing S^{-1} or any inverse, and points out new easy-to-apply preconditioners that exploit the linear invariances of the original objective function. We illustrate the consequences of the approach on optimization problems involving ReLU networks and their rescaling symmetries [3, 4].

References

[1] Boumal, Nicolas, An introduction to optimization on smooth manifolds, Cambridge University Press (2023)

[2] Ollivier, Yann, Riemannian metrics for neural networks I: feedforward networks, Information and Inference: A Journal of the IMA (2015)

[3] Marco Armenta and Thierry Judge and Nathan Painchaud and Youssef Skandarani and Carl Lemaire and Gabriel Gibeau Sanchez and Philippe Spino and Pierre-Marc Jodoin, Neural Teleportation, Mathematics (2023)

[4] Bo Zhao and Nima Dehmamy and Robin Walters and Rose Yu, Symmetry Teleportation for Accelerated Optimization, Advances in Neural Information Processing Systems (2022)

Session 10B: Thoughts on multi-objective optimization *Room: Lise Meitner (A1.116)*

What is the Objective? Investigating Design Decisions in Continuous Multi-Objective Optimization

Diederick Vermetten

Sorbonne Université, CNRS, LIP6
diederick.vermetten@lip6.fr

Keywords: Multi-Objective Optimization, Algorithmic Bias, Performance Evaluation, Restarts

The structure of an optimization problem determines what methods can be used to solve it. One key distinction arises immediately at problem formulation: is there a single objective function, or are there multiple? Since this choice determines whether a complete ordering exists between potential solutions, it influences all subsequent algorithm design choices. As a result, multi-objective optimization is often analyzed as its own subdomain, with its own set of practices, which can diverge from those present in the single-objective setting. One key difference between the settings lies in how we analyze the performance of an algorithm. When dealing with multiple objectives, we need to judge not just a single point, but rather a set of tradeoff solutions. This means that concepts like a best-so-far solution require some changes to translate. In the talk, we will showcase how we tackle this in the IOHinspector framework [2], making use of an unbounded archive to regain monotonicity for anytime analysis. Next, we will discuss how the design of benchmark problems, with a focus on the structure of the objective space, can lead to potential bias in the way multi-objective optimization algorithms are designed. By placing the focus only on the objective space, the characteristics of the search space risk being ignored, resulting in characteristics which can (unintentionally) be exploited by common, but biased, mutation operators [1]. We will finally zoom in on algorithm restarts. While these are considered natural in singleobjective optimization, most multi-objective optimizers opt to exclude them. This can in part be explained by the more exploratory focus induced by the tradeoff solutions present in the population. However, through a set of experiments we illustrate that even these methods can benefit from independent restarts when faced with sufficiently challenging multi-modal problems.

References

[1] D. Vermetten and J. Rook. Exploration of pareto-preserving search space transformations in multi-objective test functions. In Proceedings of the Genetic and Evolutionary Computation Conference, pages 50–58, 2026.

[2] D. Vermetten, J. Rook, O. L. Preuß, J. de Nobel, C. Doerr, M. López-Ibañez, H. Trautmann, and T. Bäck. Mo-iohinspector: Anytime benchmarking of multi-objective algorithms using iohprofiler. In International Conference on Evolutionary Multi-Criterion Optimization, pages 242–256. Springer, 2025.

Session 5F: IS : Recent Advances in Systematic and Computer-Aided Methods in Optimization I *Room: Ada Lovelace (A1.128)*

Can Acceleration in Gradient-Norm Minimization Be Anytime? Sharp Last-Iterate Limits in Smooth Convex Optimization

Pierre Vernimmen[1], François Glineur[1]

[1]UCLouvain, Belgium, pierre.vernimmen@uclouvain.be, françois.glineur@uclouvain.be

Keywords: Gradient Norm, Anytime Algorithms, First-Order Methods, Lower Bounds

In smooth convex optimization, the gradient norm is a directly observable measure of stationarity. Accelerating a first-order method that minimizes the gradient norm is known to be more delicate than accelerating the minimization of function values. Optimal accelerated methods such as OGM-G [1] exist for any prescribed finite horizon, but their coefficients depend explicitly on that horizon, i.e. on the number of iterations. We ask what kind of acceleration remains possible when the stopping horizon is unknown to the method, i.e. for horizon-independent methods. Diakonikolas and Wang [2] conjectured that, for any nonadaptive horizon-independent linear-span first-order method, an $\Omega(N^{-1})$ lower bound on the squared gradient norm must hold at every horizon N. Recent work [3] proves such a pointwise obstruction for the more restricted class of nonadaptive gradient descent with positive stepsizes, supporting this conjecture. We show, however, that the conjecture is false for the broader linear-span class: a single horizon-independent method can achieve near-N^{-2} last-iterate guarantees on a density-one set of horizons. What remains true is that an $\Omega(N^{-1})$ lower bound must occur infinitely often. More precisely, if $G_N(A)$ denotes the bound on the squared gradient norm after N iterations for a method A, we prove that any method in the class satisfies $\limsup_{N\to\infty} NG_N(A) \geq 1/2$. This bound is sharp: the constant $1/2$ is exactly attained by the horizon-independent gradient-descent schedule of [4]. We further show that the two extreme behaviors cannot coexist for the same method: if the bound for method A satisfies $G_{N_j}(A) = o(N_j^{-1})$ along some subsequence, then we also necessarily have $\limsup_N NG_N(A) = \infty$. Thus, horizon independence does not rule out acceleration at almost every horizon, but it prevents uniformly accelerated last iterates. In contrast, best-so-far output admits a uniform $O(N^{-2})$ guarantee.

References

[1] D. Kim and J. A. Fessler. Optimizing the efficiency of first-order methods for decreasing the gradient of smooth convex functions. J. Optim. Theory Appl., 188(1):192–219, 2021.

[2] J. Diakonikolas and P. Wang. Potential function-based framework for minimizing gradients in convex and min-max optimization. SIAM J. Optim., 32(3):1668–1697, 2022.

[3] C.-E. Tsai, I. Fatkhullin, L. Zhang, and N. He. Lower bounds for anytime acceleration of gradient descent. arXiv preprint arXiv:2607.02053, 2026.

[4] T. Rotaru, F. Glineur, and P. Patrinos. Exact worst-case convergence rates of gradient descent: A complete analysis for all constant stepsizes over nonconvex and convex functions. Math. Program., 2026. doi:10.1007/s10107-025-02313-1.

Session 12I: Optimization for Learning and Large-Scale Decision Making *Room: Sophie Germain (A1.139)*

Primal-Dual Coordinate Descent for Nonconvex-Nonconcave Saddle Point Problems Under the Weak MVI Assumption

Iyad Walwil[1], Olivier Fercoq[2]

[1]An-Najah National University, Palestine, iyad.walwil@najah.edu
[2]Télécom Paris, France, olivier.fercoq@telecom-paris.fr

Keywords: Nonconvex optimization, Nonconvex proximal, Weak MVI, Coordinate descent, Random extrapolation, Computer-aided proofs.

We introduce two primal-dual algorithms for nonconvex-nonconcave nonsmooth saddle point problems (NC-SPP) characterized by the weak Minty Variational Inequality (MVI) [3]. The first, Nonconvex-Nonconcave Primal-Dual Hybrid Gradient (NC-PDHG), extends the PrimalDual Hybrid Gradient (PDHG) method [2]. The second, Nonconvex-Nonconcave Stochastic Primal-Dual Hybrid Gradient (NC-SPDHG), incorporates a randomly extrapolated primal-dual coordinate descent approach, extending the Stochastic PDHG (SPDHG) algorithm [1]. To our knowledge, designing a coordinate-based algorithm for NC-SPP is unprecedented, and proving its convergence posed significant difficulties. This challenge motivated us to utilize PEPit [5], a Python-based tool for computer-assisted worst-case analysis of first-order optimization methods. By integrating PEPit with automated Lyapunov function techniques [4], we successfully derived the NC-SPDHG algorithm. Both methods converge under a mild condition on the weak MVI parameter, using constant, problem-adapted step sizes. Experiments on sigmoid regression with squared loss and perceptron regression confirm our theory and outperform state-of-theart methods, where linear convergence is observed. Additionally, we conduct a convex-concave least-squares experiment to show that NC-SPDHG performs competitively with SAGA, a leading algorithm in the smooth convex setting.

References

[1] A. Chambolle, M. J. Ehrhardt, P. Richtárik, and C.-B. Schönlieb. Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. SIAM Journal on Optimization, 28(4):2783–2808, 2018.

[2] A. Chambolle and T. Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. Journal of Mathematical Imaging and Vision, 40(1):120–145, May 2011.

[3] J. Diakonikolas, C. Daskalakis, and M. Jordan. Efficient methods for structured nonconvexnonconcave min-max optimization. In A. Banerjee and K. Fukumizu, editors, Proceedings of The 24th International Conference on Artificial Intelligence and Statistics, volume 130 of Proceedings of Machine Learning Research, pages 2746–2754. PMLR, 13–15 Apr 2021.

[4] O. Fercoq. Defining lyapunov functions as the solution of a performance estimation saddle point problem, 2024.

[5] B. Goujaud, C. Moucer, F. Glineur, J. M. Hendrickx, A. B. Taylor, and A. Dieuleveut. Pepit: computer-assisted worst-case analyses of first-order optimization methods in python. Mathematical Programming Computation (to appear), 2024.

Session 8G: Optimizing under network structure *Room: Enrico Fermi (A1.133)*

Online Design of Dynamic Networks

Duo Wang, Andrea Araldo, Mounîm El-Yacoubi

1Peking University, China 2Institut Polytechnique de Paris, France, araldo@telecom-sudparis.eu

Keywords: Network design, stochastic control, Monte Carlo tree search, imitation learning, public transport

Network design is classically performed offline, before the network starts operating. This is restrictive when the environment is stochastic and evolves during operation: a transport network, a communication network, or a service system may benefit from adapting its topology to requests and costs that are only progressively revealed. We introduce the problem of online design of dynamic networks: the network itself is a decision variable that evolves over time, rather than a fixed infrastructure over which online routing decisions are subsequently made [1]. We represent the evolving network as a time-expanded graph built over a fixed substrate graph. A stochastic environment generates requests that have to be routed on the graph being constructed. Network modifications are formulated as an impulsive control strategy: at each intervention, the controller selects which time-expanded edge to add and when it should become active. Under natural assumptions, we prove that optimal interventions can be restricted to specific discrete decision epochs. This reduces the continuous-time control problem to a Markov decision process, while retaining the temporal structure needed to route requests over the evolving network. The resulting state-action space is nevertheless combinatorial: every decision opens many possible future construction trajectories. Moreover, some parts of the network must be committed in advance so that users or routing algorithms can plan future paths. We address these two difficulties with a rolling-horizon method based on Monte Carlo Tree Search (MCTS). Search is guided by a neural-network policy learned by imitation from a limited set of exemplary network evolutions, while a prediction model simulates future environmental dynamics. The learned policy guides MCTS toward promising network structures, while MCTS still determines the online control decisions. We illustrate the approach on dynamic public transport. Instead of extending each vehicle trajectory independently, as in dynamic vehicle-routing methods [2], the controller constructs a structured network of bus lines on the fly, allowing passenger journeys that include transfers between lines. Experiments driven by real New York City trip requests show that this network-wide structure nearly doubles the service rate relative to a state-of-the-art high-capacity ride-sharing method under comparable operating conditions. Comparisons with static bus-network design also show the value of adapting the network online. Additional case studies on evolvable complex systems and the k-server problem demonstrate that the formulation is not transport-specific. The work thus connects stochastic optimal control, combinatorial search, and learning to make network topology an online operational decision.

References

[1] D. Wang, A. Araldo, M. El-Yacoubi, “Online Design of Dynamic Networks,” to appear in Proc. IEEE Conference on Decision and Control, 2026. Extended version: arXiv:2410.08875. [2] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, D. Rus, “On-demand high-capacity ridesharing via dynamic trip-vehicle assignment,” Proc. Natl. Acad. Sci. USA, 114(3):462–467, 2017.

Session 9A: Physics Informed Neural Networks *Room: Marie Curie (Auditorium)*

Learning PDEs for Portfolio Optimization with Quantum Physics-Informed Neural Networks

Letao Wang, Abdel Lisser, Sreejith Sreekumar, Zeno Toffano

1Laboratory of Signals and Systems, CentraleSupélec, CNRS, Paris-Saclay University, France 2Fédération de Mathématiques de CentraleSupélec, CentraleSupélec, CNRS, Paris-Saclay University, France
letao.wang@centralesupelec.fr, abdel.lisser@centralesupelec.fr sreejith.sreekumar@centralesupelec.fr, zeno.toffano@centralesupelec.fr

Keywords: Quantum machine learning, Variational quantum algorithm, Quantum PhysicsInformed Neural Network, Portfolio optimization, Partial differential equation.

Partial differential equations (PDEs) play a crucial role in financial mathematics, particularly in portfolio optimization, and solving them using classical numerical or neural network methods has always posed significant challenges. Here, we investigate the potential role of quantum circuits for solving PDEs. We design a parameterized quantum circuit (PQC) for implementing a polynomial based on tensor rank decomposition, reducing the quantum resource complexity from exponential to polynomial when the corresponding tensor rank is moderate. Building on this circuit, we develop a Quantum Physics-Informed Neural Network (QPINN) [1] and a Quantuminspired PINN, both of which guarantee the existence of an approximation of the PDE solution, and this approximation can be represented as a polynomial that incorporates tensor rank decomposition. Numerical experiments are conducted on the Hamilton–Jacobi–Bellman (HJB) PDE arising from the Merton portfolio optimization problem [2], which determines the optimal investment fraction between a risky and a risk-free asset. Unlike static quantum portfolio optimization [3], our QPINN approximates a continuous-time value function, enabling repeated evaluation and strategy recovery from its derivatives without retraining. The results show that our quantum models achieve lower losses and approximation errors than a classical fully connected PINN while using substantially fewer trainable parameters. Our quantum models further outperform a classical PINN constructed to share a similar inductive bias, providing experimental evidence of quantum-induced improvement in the tested settings and highlighting a resourceefficient pathway toward classical and near-term quantum solvers for PDEs with exploitable solution structure.

References

[1] O. Kyriienko, A. E. Paine, V. E. Elfving, Solving nonlinear differential equations with differentiable quantum circuits, Physical Review A, 103:052416, 2021.

[2] R. C. Merton, Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, The Review of Economics and Statistics, 51(3):247–257, 1969. doi:10.2307/1926560.

[3] P. Rebentrost, S. Lloyd, Quantum Computational Finance: Quantum Algorithm for Portfolio Optimization, KI – Künstliche Intelligenz, 38:327–338, 2024.

Session 12H: Market Design and Strategic Behavior *Room: Alan Turing (A1.134)*

Long-time behavior of optimal particle systems in mean-field control

Songbo Wang[1], Wenbin Yan[2]

[1]Université Paris Dauphine-PSL, France, wang@ceremade.dauphine.fr
[2]Université Paris Dauphine-PSL, France, wenbin.yan@dauphine.eu

Keywords: mean-field control, propagation of chaos, turnpike phenomenon, modulated free energy

We provide quantitative estimates on the long-time behavior of the optimal particle system associated with a mean-field control problem. Our analysis covers both displacement and flat convex settings, and extends to flat semiconvex regimes in which the concave part is absorbed by heat dissipation. Combining the modulated free energy method with functional inequalities on N-particle configurations, we compare arbitrary particle flows to their mean-field counterpart, obtaining a uniform-in-time propagation of chaos, as well as a turnpike estimate around the mean-field ergodic state up to an approximation error.

Session 8D: Invited session: Recent Advances in Systematic and Computer-Aided Methods in Optimization part 2
Room: Lise Meitner (A1.116)

Optimized Projection-Free Algorithms for Online Learning: Construction and Worst-Case Analysis

Julien Weibel, Pierre Gaillard, Wouter Koolen, Adrien Taylor

[1]Inria, D.I. École Normale Supérieure, CNRS, PSL Research University, France; [2]Université Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France; [3]CWI, Amsterdam and Twente University, Netherlands

Keywords: online Frank–Wolfe, algorithm design, online convex optimization, Performance Estimation Problem, Semidefinite programming

This work studies and develops projection-free algorithms for online learning with linear optimization oracles (a.k.a. Frank–Wolfe) for handling the constraint set, and for convex loss functions. More precisely, this work (i) shows how to exploit semidefinite programming to jointly design and analyze online Frank–Wolfe-type algorithms numerically in a variety of settings, (ii) leverages those design techniques to propose an improved (optimized) variant of an online Frank–Wolfe algorithm along with its conceptually simple potential-based proof, and (iii) extends this proof to its anytime version, which benefits from a similar $O(T^{3/4})$ regret rate without requiring knowledge of the time horizon T in advance. We are not aware of other direct regret guarantees for an anytime version of online Frank–Wolfe without using the classical doubling trick. Based on the semidefinite technique, we conclude with strong numerical evidence suggesting that no pure online Frank–Wolfe algorithm within our model class can have a regret guarantee better than $O(T^{3/4})$ without additional assumptions, that the current algorithms do not have optimal constants, and that multiple linear optimization rounds do not generally help to obtain better regret bounds.

References

[1] J. Weibel, P. Gaillard, W. Koolen, A. Taylor, Optimized Projection-Free Algorithms for Online Learning: Construction and Worst-Case Analysis, Proceedings of The 29th Interna- tional Conference on Artificial Intelligence and Statistics (AISTATS 2026), 300:5122-5130, 2026.

Session 5J: Invited session : Quantum Applications in Operations Research part 1 *Room: Emmy Noether (A1.140)*

Quantum Approaches to Multi-Objective Optimization

Laura Wendel, Alexey Bochkarev, Zakaria Abdelmoiz Dahi, Pascal Halffmann

[1]RPTU Kaiserslautern-Landau, Germany; [2]Fraunhofer ITWM, Germany; [3]Inria, France

Keywords: Quantum computing, multi-objective optimization, discrete optimization

This talk discusses quantum approaches to multi-objective optimization. Current literature on the topic is relatively sparse and can be loosely structured around three lines of research. The first one focuses on sampling the solutions of scalarized single-objective reformulations of the problem. From the quantum computing perspective, such approaches include both gate- based (with [5] being a prominent example) and annealing-based algorithms [1, 4]. The second line of work adapts variational algorithms, most notably Quantum Approximate Optimization Algorithm (QAOA), to multi-objective setting, for example, [3, 6]. Finally, there are approaches relying on Grover's search subroutine, such as [2]. We outline the key algorithms present in the literature and discuss their advantages and limitations. To highlight the resulting computational dynamics, we consider a few numerical examples of multi-objective problem instances and discuss possible approaches to sampling locally- and globally-nondominated solutions.

References

[1] Benjamín Barán and Marcos Villagra. Multiobjective Optimization in a Quantum Adiabatic Computer. Electronic Notes in Theoretical Computer Science, 329:27–38, December 2016.

[2] Benjamín Barán and Marcos Villagra. Multiobjective Optimization Grover Adaptive Search. In: Stefka Fidanova (eds) Recent Advances in Computational Optimization. Studies in Computational Intelligence, vol 795, Springer, Cham, September 2018.

[3] Linus Ekstrom, Hao Wang, and Sebastian Schmitt. Variational Quantum Multiobjective Optimization. Physical Review Research, 7(2):023141, May 2025.

[4] Andrew D. King. Multi-Objective Optimization by Quantum Annealing, November 2025. Preprint, arXiv:2511.01762.

[5] Ayse Kotil, Elijah Pelofske, Stephanie Riedmüller, Daniel J. Egger, Stephan Eidenbenz, Thorsten Koch, and Stefan Woerner. Quantum Approximate Multi-Objective Optimization. Nature Computational Science, pp 1–10, October 2025.

[6] Ivica Turkalj, Tom Ewen, Pascal Halffmann, Janik Maciejewski, Michael Trebing, and Za- karia Abdelmoiz Dahi. Enhancing Variational Quantum Algorithms for Multicriteria Op- timization. Preprint, arXiv:2506.22159, June 2025.

Session 5G: Invited Session : Optimal trading in intraday electricity markets *Room: Enrico Fermi (A1.133)*

Opportunities in Algorithmic Trading for Battery Energy Storage Systems in Short-Term Electricity Markets

David Wozabal

Vienna University of Economics and Business, Austria

Keywords: Algorithmic trading, battery energy storage, intraday electricity markets, ancillary services, high-frequency trading

This talk explores the potential, strategies, and challenges of high-frequency algorithmic trading in continuous intraday power markets. We examine key market characteristics and explore the design of profitable high-frequency trading strategies. Additionally, we discuss the opportunities for multi-market trading, integrating the intraday market with ancillary service markets.

Session 15C: Energy Systems *Room: André Marie Ampère (Amphi II)*

Profit-Optimized Maintenance Task Clustering for Gas-Insulated Substations

Wenbo Wu[1], Rim Kaddah[1], Miguel F. Anjos[2], Sihem Tebbani[3]

[1]Technological Research Institute SystemX, Palaiseau, France
{wenbo.wu, rim.kaddah}@irt-systemx.fr
[2]School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh EH9 3FD, and GERAD HEC Montreal, Montreal (Quebec), Canada H3T 2A7.
anjos@stanfordalumni.org
[3]Université Paris-Saclay, CentraleSupélec, CNRS, Laboratoire des Signaux et Systèmes, France
sihem.tebbani@centralesupelec.fr

Keywords: Predictive maintenance, Task grouping, integer linear programming, cost optimization, gas-insulated substations

Gas-insulated substations (GIS) are critical components of power transmission systems and require recurrent sulfur hexafluoride refilling to preserve dielectric integrity. Scheduling maintenance interventions independently at their predicted execution dates can lead to fragmented plans, repeated crew dispatches, and high coordination costs. We propose a profit-maximization framework for consolidating predictive maintenance tasks across GIS compartments. The approach combines physics-informed pressure degradation models with an integer linear programming formulation, allowing tasks to be shifted within pressure-based safety windows subject to daily capacity limits. The objective balances dispatch-cost savings from task consolidation against asymmetric time-shift penalties, with delays penalized more heavily than advances. The computational study compares three strategies: M1, a cluster-minimization strategy developed in our previous work [1]; M2, the proposed profit-maximization model; and M3, a fixedday greedy heuristic. A sliding-window procedure extends the optimization over long planning horizons while preserving feasible task spacing. Results show that M1 reduces the number of maintenance groups by 30–60% but incurs substantial time-shift penalties, whereas M2 achieves the lowest total cost, reducing it by 50–70% relative to M1 and by 12–22% relative to M3. M3 offers competitive cost performance with lower computational complexity but requires more maintenance groups. Sensitivity analysis identifies setup-dominated, penalty-constrained, and saturation regimes, highlighting the trade-off between task consolidation, scheduling flexibility, and operational cost.

References

[1] W. Wu, R. Kaddah, M. Anjos, and P.-H. Langlois, "Task grouping optimization for gas insulated substations (GIS) predictive maintenance: Trade-off between cost and environmental impact," in Proceedings of the 13th IMA International Conference on Modelling in Industrial Maintenance and Reliability, July 2025. doi: 10.19124/ima.2025.01.08.

Session 8E: Optimisation with chance constraints *Room: Antoine Lavoisier (A1.122)*

Eventual convexity for separable chance constraints with skewed generalized hyperbolic random variables

Heng Zhang, Abdel Lisser

1Université Paris-Saclay, CNRS, CentraleSupelec, Laboratory of Signals and Systems, France, heng.zhang@centralesupelec.fr 2Université Paris-Saclay, CNRS, CentraleSupelec, Laboratory of Signals and Systems, France, abdel.lisser@centralesupelec.fr

Keywords: Chance constraints, Skewed distributions, Generalized hyperbolic distributions, Eventual convexity

Chance constraints are widely used in optimization under uncertainty. This paper aims to show the eventual convexity (EV) of chance constraints with skewed generalized hyperbolic (GH) random variables. We prove that the densities of GH distributions are α-decreasing, and obtain exact convex reformulations for separable jointly chance constraints with GH distributions. We provide numerical results to compute the α-decreasing threshold parameters and show the EV of the feasible set together with the computational tractability to solve the associated optimization problems.

Session 15I: Geometry *Room: Sophie Germain (A1.139)*

Second-Order Reliability Method based on Gram-Charlier A series with analytical Gradient and Hessian of failure probability

Qian Zhang, Mårten Olsson, Rami Mansour

1Department of Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, Sweden qiazhang@kth.se, mart@kth.se 2Department of Mechanical and Production Engineering, Aarhus University, Aarhus, Denmark ramimansour@mpe.au.dk

Keywords: Probability of failure, Gradient of failure probability, Hessian of failure probability, Most Probable Point of failure, Second-Order Reliability Method

Reliability-based design optimization requires accurate first- and second-order derivatives of the failure probability to construct the gradient and Hessian of the Lagrangian. For computationally expensive implicit limit-state models, finite-difference estimation is often impractical because each perturbation requires an additional reliability analysis. Moreover, first-order reliability approximations may be insufficient for strongly nonlinear limit states [1]. This work derives analytical gradient and Hessian of the failure probability within a second-order reliability framework based on the Gram-Charlier A series. All sensitivities are obtained from a single search for the most probable failure point (MPP). For implicit limit-state functions, a quasi-Newton approximation of the limit-state Hessian is recovered from the search trajectory without additional model evaluations. Numerical examples demonstrate close agreement with Monte Carlo (MC) [2] reference results and a substantial reduction in computational effort compared with second-order reliability analysis using finite-difference sensitivities.

References

[1] Z. Hu, R. Mansour, M. Olsson, and X. Du, Second-order reliability methods: a review and comparative study, Structural and Multidisciplinary Optimization, 64(6):3233–3263, 2021. doi:10.1007/s00158-021-03013-y.

[2] S.-K. Au, Optimality conditions for MCMC in rare event risk analysis, Reliability Engineering & System Safety, 265:111539, 2025. doi:10.1016/j.ress.2025.111539.

Session 15K: Multiobjective Optimization and Evolutionary Computation Part II *Room: Loggia (A1 - first floor)*

A Crowding Distance That Provably Solves the Difficulties of the NSGA-II in Many-Objective Optimization

Weijie Zheng, Yan Gao, Benjamin Doerr

1School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China 2LIX, École Polytechnique, IP Paris, Palaiseau, France, firstname.lastname@polytechnique.edu

Keywords: Multi-objective optimization, evolutionary algorithms, runtime analysis.

Recent mathematical runtime analyses have shown that the crowding distance used by the NSGA-II multi-objective optimizer does not well reflect how isolated, and hence worth keeping, solutions are. This has provably led to drastic difficulties when the number of objectives is three or more. Based on the insights of these works, we design a variant of the crowding distance, called truthful crowding distance, and analyze the resulting truthful NSGA-II with mathematical means. (i) We show that it solves the many-objective versions of the OMM, COCZ, LOTZ, and OJZJk problems in the same (polynomial) asymptotic runtimes as the NSGA-III and the SMS-EMOA, contrasting the exponential lower bounds for the classic NSGA-II. (ii) For the biobjective versions of these problems, our NSGA-II maintains the good performance of the classic NSGA-II, but is successful already with population sizes smaller by a constant factor. (iii) For the bi-objective OMM problem, we observe a (minimally) better performance in approximating the Pareto front. These results indicate that our truthful crowding distance is an interesting alternative to the classic crowding distance. This work is part of the PGMO-funded project From Mathematical Runtime Analyses to Better Evolutionary Multi-Objective Optimizers (PI: Benjamin Doerr).

References

[1] W. Zheng, Y. Gao, and B. Doerr. A crowding distance that provably solves the difficulties of the NSGA-II in many-objective optimization. IEEE Transactions on Evolutionary Computation, 2025. To appear. Early access version at https://doi.org/10.1109/TEVC.2025.3638306.

Author Index

Ali Abbassi, 7
Wim Van Ackooij, 45
Wim van Ackooij, 8, 9, 27, 56, 83
Steeven B. Affognon, 10
Andrei Agrachev, 11
Andrei A. Agrachev, 12
Hamid Ben Ahmed, 44
Asmae Alami, 13
Emanuel Aldea, 121
Xavier Allamigeon, 14
R. Andrade, 51
Ekaterina ANDREEVA, 15
Panagiotis Andrianesis, 16
Miguel F. Anjos, 17, 62, 146, 169
Denis Antipov, 18
Andrea Araldo, 19, 63, 93, 143, 163
Ayse N. Arslan, 20
Christian Artigues, 99
Hélène Arvis, 21
Pierre Arvy, 22
Galit Ashkenazi-Golan, 23
Laiba Asmatullah, 24
Pierre-Cyril Aubin-Frankowski, 25
Gennaro Auricchio, 92
Didier Aussel, 26
Shadi Sharif Azadeh, 105
Mathis Azema, 27

Milos Balac, 115
Guillaume Ballif, 28
Boris Baros, 29
Olivier Baude, 20
Antonin Bavoil, 30
Cristina Bazgan, 31
Olivier Beaude, 21
Pascale Bendotti, 32, 99
Inès Benito, 33
Lucas Berterottière, 34, 69
Nicolas Bessin, 35
Pascal Bianchi, 36
Anne Blavette, 86, 121
Alexey Bochkarev, 167
Olivier Bokanowski, 37, 38
Jérôme Bolte, 39, 40
Martin Bombardelli, 41
Benoît Bonnet-Weill, 42
Mustapha Bouhtou, 84
Marie Bouilloud, 43
Akram Boukabache, 129
Eric Bourreau, 102
Virginie Boutueil, 74
Louis Bouvier, 60
Claire Boyer, 123
Matias Bresciano, 44
Luce Brotcorne, 45, 62, 105
Amir Brudner, 19
Luca Brunod Indrigo, 5
Christoph Buchheim, 1
Thomas Buchholtzer, 96
Abdellah Bulaich Mehamdi, 45

François Caffier, 120
Enrico Calandrini, 46, 47
Guy Camilleri, 86, 121
Justin Capgras, 48, 49
Florin Capitanescu, 22
J.-G. Caputo, 51
Coralia Cartis, 59
Margarida Carvalho, 139
Nathan De Carvalho, 50
G. Cavalcante, 51
Bastien Cavarretta, 52
Domenico Mergoni Cecchelli, 23
Youssef Ouazzani Chahdi, 50
Tijani Chahed, 63
Haitam Chakori, 48, 49
Meynard Charles, 53
Sandrine Charousset, 10
Jean-François Chassagneux, 38, 54
Pierre Chaufour, 106
Vasileios Cheimonidis, 16
Yiyuan Chen, 14
Carolina Chiu, 8
Yu-Cheng Chiu, 55
Morgan Chopin, 31
Alexandre Chotard, 128
Yi Chu, 56
Evgenii Chzhen, 57
Samuel Cohen, 29
Samuel N. Cohen, 58
Valentin Colliard, 24
Daniel Cortild, 59, 103
Nicolas Corvol, 60
Romain Cosson, 6
Luciano Costa, 125
Axelle Cousinou, 61
Juan J. Cuenca, 121

Zakaria Abdelmoiz Dahi, 167
Guillaume Dalle, 35, 142
Claudia D'Ambrosio, 15, 48, 49, 85, 109
Nguyen Dang, 127
Ezra Daniels, 62
Mandar Datar, 63
Alexis David, 64

Andrea Davini, 65
Nadjette Debbouche, 129
Solène Delannoy-Pavy, 66
François Delarue, 67
Sebastien Delbos, 61
Lorenzo Demaria, 68
Sophie Demassey, 106, 141
Natello Descormier, 69
Boris Detienne, 20
A. Dieuleveut, 151
Aymeric Dieuleveut, 152
Sébastien Le Digabel, 146
Mao Fabrice Djete, 70
Benjamin Doerr, 33, 71, 72, 80, 172
Carola Doerr, 18, 127
Enrico Le Donne, 150
Radu-Alexandru Dragomir, 36
Yann Dujardin, 7
Alexandre Dupont-Bouillard, 73

Prince Edorh, 77
Rose Sossou Edou, 74
Rachid El-Azouzi, 135
Mounîm El-Yacoubi, 93, 163
Sourour Elloumi, 108
Moaad Elmoutassim, 71
Ikhlas Enaieh, 75
Faniriana Rakoto Endor, 76
Hamza Ennaji, 145
Carlos Esteve-Yagüe, 37, 154
Laura Bernardo Estévez, 77

Michele Fabi, 48, 49, 78
Jalal Fadili, 133, 145
Filippo de Feo, 58
Raphaël Féraud, 86
Olivier Fercoq, 75, 158, 162
Adeline Fermanian, 34
Luca Ferrarini, 92
Émilien Flayac, 112
Gerard Fleury, 79
Pierre Fouilhoux, 69
Michel Fournié, 112
Jules Fraizier, 80
Adrien Le Franc, 64, 81
Mathieu Francois, 61
Antonio Frangioni, 46, 47
Christine Fricker, 28, 138
Aurélien Froger, 20

Pierre Gaillard, 166
Jie Gao, 93
Yan Gao, 172
Guillaume Garrigos, 103
Charles Garrisi, 82
Nicolas Gast, 21
Stephane Gaubert, 45
Stéphane Gaubert, 14, 21
Bruno Gaujal, 21
Caroline Geiersbach, 83
Garance Gérard, 84
Juliette Gerbaux, 44
Carlo A. M. Gervasi, 85
Bissan Ghaddar, 2, 85
Eymeric Giabicani, 86
Céline Gicquel, 134
Oskar Girardin, 128
François Glineur, 87, 161
Xavier Goaoc, 14
Sebastiano Nicolussi Golo, 150
Salvish Goomanee, 88
Luca Gorini, 89
Omar Goudiaby, 10
Emilien Gouffault, 90
Baptiste Goujaud, 152
Eric Gourdin, 7
Florentin Goyens, 91
Camille Grange, 92
Remi Gribonval, 113
Rémi Gribonval, 159
Rodolphe Griset, 20
Lars Grüne, 3
Fabrice Guillemin, 119
Zihao Guo, 93
Dhruv Gupta, 94
Aldo Gutiérrez-Concha, 95
Julien Guyon, 96

Pascal Halffmann, 167
Cyrian Hallermeyer, 97
Andi Han, 88
Nikolaus Hansen, 128
Yezekael Hayel, 139
Jackson Hebner, 58
Alexandre Heintzmann, 32
Didier Henrion, 112
René Henrion, 98
Bruno Hérissé, 77
Alain Hertz, 114
Massil Hihat, 35, 69
Kea Horvath, 99
Olivier Hudry, 100

Andrea Iannelli, 124
Shokhrukh Ibragimov, 101
Franck Iutzeler, 4

Quentin Jacquet, 45, 130
Fabian Jakob, 124
Jean-Marc Janin, 44
Jauny, 26
Pratik Jawanpuria, 88

Frédéric Jean, 77
Xia Jiang, 93
Tania Jiménez, 139
Natalia Jorquera-Bravo, 9

Rim Kaddah, 169
Axel Kärrolm, 68
Adil Kassaoui, 102
Marc Kaufmann, 55
Sébastien Kawka, 89
Bettina Kazandjian, 11
Cheikh M. F. Kebe, 10
Safia Kedad-Sidhoum, 34, 69, 84, 108, 134
Lucas Ketels, 103
Rashed Khanjani-Shiraz, 104
Guy Mushiata Kibulu, 71
A. Knippel, 51
Wouter Koolen, 166
Milan Korda, 112
Mesut Can Köseoğlu, 105
Malak Al Kostit, 106
Martin S. Krejca, 71, 72, 80
Jan Kronqvist, 68

Safwan Labbi, 107
Philippe Lacomme, 7, 41, 79
Marguerite Lafeuillade, 108
Kostas Lakis, 55
Vinith Lakshmanan, 13
Amélie Lambert, 22, 84
Michael Laporte, 109
Michel De Lara, 110
Roman Le Goff Latimier, 44
Pierre Lavigne, 111
Alexey Lazarev, 148
Charlie Lebarbé, 112
Arthur Lebeurrier, 113, 159
Vincent Leclère, 27, 64, 66
Benoît Legat, 130
Aymeric Legros, 114
Baptiste Leloup, 78
Phong Le, 127
Quoc-Tung Le, 40
Baptiste Leroux, 115
Antoine Lesage-Landry, 146
Tam Le, 116
Quentin Lété, 130
Leo Liberti, 48, 49, 73
Tanguy Lindegaard, 117
Abdel Lisser, 118, 164, 170
Emily Little, 90
Johannes F. Lutzeyer, 33

Raneem Madani, 118
Ahlam MAGNOUCHE, 15
Patrick Maillé, 86
Paul Malisani, 119, 120
Paul Mangold, 107
Rami Mansour, 171
Alban Marie, 121
Alexandre Marié, 90
Mattia Martini, 70
Rui Ma, 122
Rodrigo Maulen-Soto, 123
Alessa Mayer, 19
L. Melo, 51
Pierre Mendy, 10
Govind Menon, 157
Donato Meoli, 47
Panayotis Mertikopoulos, 23
Alizée Mesnard, 24
Frédéric Meunier, 35, 74, 109
Jared Miller, 124
Bamdev Mishra, 88
Cecília Mota, 125
Michele Motta2., 12
Eric Moulines, 107
Antonio Mucherino, 73
Ilkhom Mukhammadiev, 101
Ravil Mussabayev, 126
Rustam Mussabayev, 126

Michel Nakhla, 43
Babacar M. Ndiaye, 10
José Neto, 17
Sandra Ulrich Ngueveu, 99
Tai Nguyen, 127
André Nichterlein, 31

Marc Oedingen, 128
Welington de Oliveira, 8, 56, 82
Mårten Olsson, 171
Nicolas Omont, 22
Stanley Osher, 154
Adel Ouannas, 129
Thomas Ouillon, 10
Younes Brahim Oumedjber, 129

François Pacaud, 82, 130
Bernardo Pagnoncelli, 134
Dimitri Papadimitriou, 131
Axel Parmentier, 60, 66
Edouard Pauwels, 39, 40
Raphaël Payen, 139
Angel Garcia Pedrero, 75
Teemu Pennanen, 132
Julien Pennec, 139
Pedro Pérez-Aros, 8
Nancy Perrot, 84
Juan Peypouquet, 59, 103
Laurent Pfeiffer, 120
Baptiste Plaquevent-Jourdain, 133

Edward Plumb, 23
Pierre-Louis Poirion, 9
Lara Pontes, 125
Daniel Porumbel, 22
Victor Priser, 36
Caroline Prodhon, 7

Franco Quezada, 134

Christoph Reisinger, 29
Anne Reulke, 135
David Rey, 136
Léa Ricard, 147
Camille Richer, 31
Rodolphe Le Riche, 122
Konstantin Riedl, 137
Tristan Rigaut, 64
Alessia Rigonat, 138
Luca Rizzi, 149
Bárbara Rodrigues, 62
Marius Roland, 9, 20, 62
Stefan Ropke, 85
Cécile Rottner, 20, 32, 99, 108
Clément W. Royer, 52, 91
Manuel Ruiz, 66

Rita Safi, 139
Adle Ben Salem, 96
Diyora Salimova, 101
Simon SANANES, 54
Mariam Sangaré, 20
Maxime Sangnier, 158
Raimundo Saona, 65
Enrico Sartor, 140
Antonio Sasaki, 141
Marco Scaratti, 38
Ulysse Schaller, 55
Simeon Schaub, 142
Sholom Schechtman, 57
Carsten Scherer, 124
Corbinian Schlosser, 148
Martin Schmidt, 98
Dieks Scholten, 72
Antonio Sciarretta, 13, 115
Frédéric Semet, 20
Ravi Seshadri, 19
Valentina Sessa, 106, 141
Yiding Shi, 93
Xu Shuangshuang, 143
Mario Sigalotti, 149
Francisco Silva, 95
Antonio Silveti-Falls, 133
Harman Singh, 144
Justin Sirignano, 58, 137
Wutao Si, 145
Christina G. Soldati, 146
Adrien Spagnol, 119
Konstantinos Spiliopoulos, 137
Sreejith Sreekumar, 164
Mathias Staudigl, 59
Austin J. Stromme, 157
Anand Subramanian, 125
Asrifa Sultana, 26
Hussein Suprême, 146
Thibault de Surrel, 52
Yue Su, 147
Grégoire Szymanski, 50

Matteo Tacchi-Bénard, 42, 148
Mathieu Tanré, 96
Bassel Tarbush, 23
A. Taylor, 151
Adrien Taylor, 152, 166
Sihem Tebbani, 169
Nicolò Tedesco, 149
Andrea Tettamanti, 150
D. B. Thomsen, 151
Daniel Berg Thomsen, 152
Daniil Tiapkin, 107
Franck Tirard, 106
Yosef Gavriel Tirat-Gefen, 153
Zeno Toffano, 118, 164
Mikaël Touati, 135
Cheik Traoré, 39
Ernesto Treumún, 155
Richard Tsai, 37, 154

Manu Upadhyaya, 152
Fernanda Urrea, 155, 156

Adrien Vacher, 157
Shivani Valecha, 26
Titouan Vayer, 113
Eyal Vayness, 158
Romain Veltz, 30
Manon Verbockhaven, 159
Diederick Vermetten, 18, 160
Pierre Vernimmen, 87, 161
Cyrille Vessaire, 66
Emilio Vilches, 95
Filipe Vinci, 153
Yannick Viossat, 23

Irène Waldspurger, 76, 91
Iyad Walwil, 162
Duo Wang, 143, 163
Letao Wang, 164
Songbo Wang, 165
Xavier Warin, 38
Julien Weibel, 166
Laura Wendel, 167
Lukas Wessels, 67

David Wozabal, 168
Wenbo Wu, 169
Yaoxin Wu, 93

Wenbin Yan, 165
Florian Yger, 52

Nicola Zaupa, 42
Heng Zhang, 170
Qian Zhang, 171
Weijie Zheng, 172
Hasnaa Zidani, 83
Bruno Ziliotto, 65
Manuel Zobel, 124
Alexandre Manso de Zuniga, 16