Third Derivative-Free Optimization Symposium

DFOS26

Université de Montréal | June 22–26, 2026

André-Aisenstadt Building, Room 6214

Scientific Committee

Charles Audet
Youssef Diouane
Warren Hare
Sébastien Le Digabel
Margherita Porcelli
Francesco Rinaldi
Stefan Wild

Contents

Derivative-Free
Optimization
Symposium

Schedule

Time	Monday 22	Tuesday 23	Wed. 24	Thursday 25	Friday 26
8:30–8:45	Registration		—	—	
8:45–9:00	Opening words	—	—	—	—
9:00–9:25	A. Lesage-Landry	Albert Berahas		Lindon Roberts	E. Hallé-Hannan
9:35–10:00	Xavier Lebeuf	Tudor Buhut	*Network*	Warren Hare	Federico Grillini
10:10–10:35	Stéphane Alarie	Katya Scheinberg	*building*	Clément Royer	Théo Denorme
10:45–11:05	*20 min break*	*20 min break*	*activities*	*20 min break*	*20 min break*
11:05–11:30	R. Filomeno Coelho	Zao Jiang		Pierre-Yves Bouchet	Matt Menickelly
11:40–12:05	Francesco Rinaldi	S. Le Digabel		Fred Gardi	Abraar Chaudhry
12:15–13:30	*lunch*	*lunch*		*lunch*	*lunch*
13:30–13:55	Scholar Sun	Liyuan Cao		Margherita Porcelli	Anthony Palmieri
14:05–14:30	Meisseme Kadri	Youssef Diouane		Christophe Tribes	Valentin Dijon
14:40–15:05	Yiwen Chen	Iman Shames		E. José da Silva	Charles Audet
15:15–15:45	*30 min break*	*30 min break*		*30 min break*	*Closing words*
15:45–16:10	Zaikun Zhang	Zhongda Huang		Nick Sahinidis	
16:20–16:45	Jeffrey Larson	Raghu Pasupathy		Paul Saves	
16:55–17:20	Stefan Wild	Zikai Xiong		Mickael Binois	

Applications	*Mo am*	Direct Search: Theory & Foundations	*Th am*
High-Dimensional DFO	*Mo pm-1*	Multiobjective Optimization	*Th pm-1*
Model-Based Methods & Surrogates	*Mo pm-2*	Data-Driven & Active Learning	*Th pm-2*
Stochastic DFO	*Tu am*	Constrained, Composite & Complex BBO	*Fr am*
Zeroth-Order & Nonsmooth Methods	*Tu pm-1*	Nonsmooth & Function-Free Methods	*Fr pm-1*
Machine Learning & Surrogates	*Tu pm-2*	Benchmarking & Software	*Fr pm-2*

Derivative-Free
Optimization
Symposium

Monday, June 22 — Morning

Time	Speaker	Title
9:00–9:25	Antoine Lesage-Landry	Blackbox optimization for electric power distribution systems
9:35–10:00	Xavier Lebeuf	Stochastic blackbox optimisation in electrical asset management with sub-sampling
10:10–10:35	Stéphane Alarie	Blackbox Optimization and Digital Twins at Hydro-Québec
11:05–11:30	Rajan Filomeno Coelho	Dimensionality Reduction based on Topological Data Analysis for the Multidisciplinary Design Optimization of a 3D Wing
11:40–12:05	Francesco Rinaldi	A Direct-Search approach for the Technology Diffusion Problem

Monday, June 22 — Afternoon

Time	Speaker	Title
13:30–13:55	Scholar Sun	Towards a Practical Geometry Correcting Trust Region Algorithm with Theory
14:05–14:30	Meisseme Kadri	A Scalable Trust-Region Framework for High-Dimensional Bayesian Optimisation
14:40–15:05	Yiwen Chen	CLARSTA: A Random Subspace Trust-region Algorithm for Convex-constrained Derivative-free Optimization
15:45–16:10	Zaikun Zhang	New perspectives on trust-region method with inaccurate gradient information
16:20–16:45	Jeffrey Larson	Composite derivative-free optimization using Birkhoff interpolation models
16:55–17:20	Stefan Wild	Hybrid Neural-Quadratic Model-Based Derivative-Free Optimization

Tuesday, June 23 — Morning

Time	Speaker	Title
9:00–9:25	Albert Berahas	Adaptive Optimization Methods and Noisy Oracles
9:35–10:00	Tudor Buhut	Chance-constrained blackbox optimization
10:10–10:35	Katya Scheinberg	Errors in stochastic derivative free optimization
11:05–11:30	Zao Jiang	An Extended Kolmogorov Inequality for Sharper Complexity Bounds in Randomized Block Coordinate Descent
11:40–12:05	Sébastien Le Digabel	A Riemannian framework for feasible DFO with equality constraints

Tuesday, June 23 — Afternoon

Time	Speaker	Title
13:30–13:55	Liyuan Cao	Complexity Analysis of Regular Simplicial Search for Minimizing Smooth Black-Box Functions
14:05–14:30	Youssef Diouane	Zeroth-Order Min-Max Optimisation for Nonconvex-Nonconcave Functions
14:40–15:05	Iman Shames	Solving the Offline and Online Min-Max Problem of Non-smooth Submodular-Concave Functions: A Zeroth-Order Approach
15:45–16:10	Zhongda Huang	Multi-Fidelity Warm-Start Optimization: Theory and Application
16:20–16:45	Raghu Pasupathy	Derivative-Free Inexact Mirror Langevin Algorithms
16:55–17:20	Zikai Xiong	Function-free Optimization via Comparison Oracles

Thursday, June 25 — Morning

Time	Speaker	Title
9:00–9:25	Lindon Roberts	Polling Set Construction and Worst-Case Complexity for Direct Search under Polyhedral Convex Constraints
9:35–10:00	Warren Hare	Practical computation of cosine measures
10:10–10:35	Clément Royer	It's a graph! It's a polytope! No, it's a PSS!
11:05–11:30	Pierre-Yves Bouchet	Convergence towards a local solution by algorithms with a covering step
11:40–12:05	Fred Gardi	Black-Box Optimization with Hexaly

Thursday, June 25 — Afternoon

Time	Speaker	Title
13:30–13:55	Margherita Porcelli	A multilevel stochastic regularized first-order method
14:05–14:30	Christophe Tribes	Blackbox multi-period power flow optimization under realistic operational constraints
14:40–15:05	Everton José da Silva	A Direct Multisearch Approach (DMS) for Many-Objective Derivative-Free Optimization
15:45–16:10	Nick Sahinidis	Data-driven optimization: efficient adaptive learning for self-driving laboratories
16:20–16:45	Paul Saves	Black-box Active Learning for Discovering New Explanations in Complex System Exploration
16:55–17:20	Mickael Binois	Multi-objective Bayesian optimization meets multiple gradient descent

Friday, June 26 — Morning

Time	Speaker	Title
9:00–9:25	Edward Hallé-Hannan	Surrogate-based categorical neighborhoods for mixed-variable blackbox optimization

Time	Speaker	Title
9:35–10:00	Federico Grillini	Boosting BFO: a new strategy to efficiently explore the integer lattice
10:10–10:35	Théo Denorme	Adaptive direct search algorithms with relaxable and quantifiable constraints
11:05–11:30	Matt Menickelly	Managing Multiple Models in Expensive Derivative-Free Optimization with Contextual Bandits
11:40–12:05	Abraar Chaudhry	On Complexity of Model-Based Derivative-Free Methods

Friday, June 26 — Afternoon

Time	Speaker	Title
13:30–13:55	Anthony Palmieri	Extrapolation-based Direct Search for Nonsmooth Stochastic Zeroth-Order Optimization
14:05–14:30	Valentin Dijon	Benchmarking Techniques for Bilevel Derivative-free Optimization
14:40–15:05	Charles Audet	Benchmarking algorithms on large test sets

Abstracts

Monday, June 22 — Morning

Blackbox optimization for electric power distribution systems

Antoine Lesage-Landry *GERAD, Mila, Polytechnique Montréal*
with Sébastien Le Digabel, Christina Soldati (GERAD, Polytechnique Montréal), Feng Li (Eaton), and Ilhan Kocar (Polytechnique Montréal)

9:00–9:25 ↑ Back to schedule

Simulators have long been used to help traditional planning and operations of electric power systems. As modern power systems become increasingly complex, relying only on seasoned engineers to interpret snapshots of the systems from these simulators becomes insufficient for decision-making. From an optimization perspective, simulators act as blackbox oracles: their internal models are inaccessible, derivatives are unavailable, and each evaluation may be computationally expensive and/or time-consuming. In this context, we design blackbox optimization methodologies to support power system operations without simplifying its physics. In this talk, we discuss two applications in distribution networks. First, we formulate an optimal reactive power dispatch stochastic program for volt-var optimization, where decision variables include both conventional device setpoints and the droop-curve of distributed energy resources. Second, we propose a network topology reconfiguration method for active power loss minimization. To improve the solution quality, algorithms inspired by variable neighbourhood search and branch-and-bound are adapted to the power-system context and combined with the blackbox optimization solver. In both applications, we consider three-phase, unbalanced distribution networks integrating sources of renewable power, illustrating how derivative-free optimization can support reliable and efficient grid operations.

- ▷ *Practical Solutions to Volt-Var Optimization Under Uncertainty via Blackbox Optimization* [32].
- ▷ *Blackbox Optimization for Loss Minimization in Power Distribution Networks using Feeder Reconfiguration* [43].

Two-Stage Blackbox Optimization with Optimal Transport-Based Sub-Sampling

Xavier Lebeuf *Polytechnique Montréal*
with Stéphane Alarie (GERAD, Hydro-Québec), Miguel Diago (Hydro-Québec), Charles Audet (GERAD, Polytechnique Montréal), Sébastien Le Digabel (GERAD, Polytechnique Montréal)

9:35–10:00 ↑ Back to schedule

This work emerges from an asset management problem at Hydro-Québec, in which a stochastic blackbox optimization (BBO) problem has prohibitive function evaluation costs. A structural property of the blackbox problem is exploited to decompose function evaluations into a fast-to-evaluate stochastic blackbox (stage 1) and an expensive deterministic blackbox (stage 2). We propose an optimal transport-based sub-sampling strategy that leverages this decomposition to obtain accurate function approximations with few calls to the stage 2 blackbox. This strategy can be integrated into various BBO solvers for stochastic problems. To test this method, we introduce MicroPRIAD, a new stochastic benchmark problem that replicates the properties of Hydro-Québec's asset management problem. Experimental results show that the proposed sub-sampling technique significantly improves optimization efficiency.

- ▷ *Risk-Based Optimization of Periodic Maintenance for Power Grid Equipment* [17].

Blackbox Optimization and Digital Twins at Hydro-Québec

Stéphane Alarie *Hydro-Québec*
with Quentin Dollon, Melad Fahed (McGill University), Arthur Favrel, Martin Gagnon, Michael Kokkolaras (Clemson University), Vincent Mai, Quang Hung Pham, Jean-Philippe Gauthier (all Hydro-Québec unless noted)

10:10–10:35 ↑ Back to schedule

Digital Twins (DTs) are virtual replicas of physical assets that enable, among other applications, anomaly detection, predictive maintenance and scenario analysis. Blackbox Optimization (BBO) is ubiquitous in DTs and is not limited to the only calibration. Hydro-Québec (HQ) is strongly involved into the development of DTs, particularly for hydropower units. This talk presents several ongoing research projects combining DTs and BBO. The calibration process will naturally be discussed [18], but also the sampling of asset behavior to reduce the epistemic uncertainty of DTs [19] as well as the fine-tuning of startup parameters, where the physical asset is directly integrated into the BBO loop and the DT is used to provide measurements as virtual sensors [33].

- ▷ *Bayesian input-state-parameter inference of hydrodynamic bearings: From partial displacement measurements to force reconstruction* [18].
- ▷ *Hydropower unit digital twin calibration using monitoring data and blackbox optimization* [19].
- ▷ *Active learning-based optimization of hydroelectric turbine startup to minimize fatigue damage* [33].

Dimensionality Reduction based on Topological Data Analysis for the Multidisciplinary Design Optimization of a 3D Wing

Rajan Filomeno Coelho *Cenaero*
with Caroline Sainvitu, Tariq Benamara (Cenaero)

11:05–11:30 ↑ Back to schedule

In the aircraft industry, the multidisciplinary design of wings is a challenging task, partly because of the high number of design variables required to parametrize in detail their geometries (skin thickness, ribs, stiffeners, etc.). To address this task, an original dimensionality reduction methodology is proposed in this study, leaning on Topological Data Analysis in order to find similarities of behaviors between variables. The methodology works as follows: after a preliminary Design of Experiments launched on numerical simulations combining Computational Fluid Dynamics (CFD) and Finite Element Method (FEM) analyses, the patterns of the responses with respect to the design variables are identified using a strategy inspired by the Uniform Manifold Approximation and Projection (UMAP) technique, initially proposed by McInnes and Healy. Variables revealed to be close in the UMAP latent space are clustered and modified concurrently during the optimization, thereby leading to a reduced design space. The global strategy is demonstrated on the aerodynamical optimization of the Onera M6 wing design.

- ▷ *Dimensionality reduction combining variable grouping and UMAP for structural optimization* [12].

A Direct-Search approach for the Technology Diffusion Problem

Francesco Rinaldi *University of Padova*
with Matteo Bergamaschi (University of Padova), Marta Pozzi, Laura Sanità (Bocconi University)

11:40–12:05 ↑ Back to schedule

The Technology Diffusion Problem (TDP) is a fundamental optimization problem on graphs, with applications ranging from innovation spread in social networks to information propagation in complex systems. Given its combinatorial nature, TDP poses significant computational challenges, especially on large-scale graphs. Currently, beyond heuristic approaches, IP formulations can provide good approximate solutions, but they often lack scalability. In this talk, we introduce two complementary contributions. First, we propose a new exact IP formulation that strengthens the theoretical modeling of the problem. Second, we present a practical direct search approach designed to efficiently tackle large-scale graphs, producing high-quality approximate solutions within reasonable computational time.

- ▷ *An efficient network-aware direct search method for influence maximization* [8].
- ▷ *The diffusion of networking technologies* [21].

Monday, June 22 — Afternoon

Towards a Practical Geometry Correcting Trust Region Algorithm with Theory

Scholar Sun *Georgia Tech*
with Abraar Chaudhry and Katya Scheinberg (Georgia Tech)

13:30–13:55 ↑ Back to schedule

Recently, a framework was proposed to analyze the complexity of a trust-region method equipped with a geometry-correcting subroutine, in the spirit of Powell's methods. We present a practical implementation that fits within this framework, providing both complexity guarantees and strong empirical performance. Among the key implementation features, we partition the sample set into a carefully managed geometry set and the remaining interpolation points. Additionally, we report preliminary findings on the algorithm's performance in a subspace setting.

▷ *On Complexity of Model-Based Derivative-Free Methods* [10].

A Scalable Trust-Region Framework for High-Dimensional Bayesian Optimisation

Meisseme Kadri *GERAD, Polytechnique Montréal*
with Amina Lamghari (Université du Québec à Trois-Rivières), Youssef Diouane, Issmaïl El Hallaoui (GERAD, Polytechnique Montréal)

14:05–14:30 ↑ Back to schedule

Bayesian optimisation (BO) is known to be very effective for solving low-dimensional expensive black-box optimisation (BBO) problems; however, in high dimensions, it often scales poorly and becomes highly sensitive to the acquisition function. To address these issues, we introduce STREGO, a scalable trust-region BO method that combines global exploration with local exploitation refinements. STREGO adapts an existing trust-region BO framework (known as TREGO) to efficiently solve high-dimensional BBO problems. Similar to TREGO, the STREGO framework is organised into two phases: a global phase and a local phase. While the TREGO global phase performs mono-objective BO, STREGO solves a mean–variance surrogate-based bi-objective problem and clusters the resulting non-dominated set to select a diverse batch of candidate points. If this global batch fails to improve the objective function, STREGO switches to a local phase, where it fits and minimises a surrogate model using only evaluations within a trust-region. Experiments on scalable analytical BBO benchmarks, as well as on classification tasks with up to 300 dimensions, indicate that STREGO outperforms not only TREGO but also state-of-the-art high-dimensional BO methods.

CLARSTA: A Random Subspace Trust-region Algorithm for Convex-constrained Derivative-free Optimization

Yiwen Chen *University of British Columbia*
with Warren Hare and Amy Wiebe (University of British Columbia)

14:40–15:05 ↑ Back to schedule

This talk proposes a random subspace trust-region algorithm for general convex-constrained derivative-free optimization (DFO) problems. As in prior random subspace DFO methods, convergence of the proposed algorithm relies on sufficient model accuracy and adequate subspace quality. To address model accuracy, we introduce a new class of models that is only required to achieve reasonable accuracy on the projection of the constraint set onto the chosen subspace. We further propose a novel geometric measure that facilitates the analysis, construction, and management of such models. To ensure subspace quality, we leverage concentration-of-measure results on the Grassmann manifold to develop a randomized subspace sampling strategy that preserves a fixed fraction of the first-order criticality measure with a prescribed probability. Building on these theoretical developments, we present an almost-sure global convergence and a worst-case complexity analysis of our algorithm. Numerical experiments on problems with dimensions up to 10000 demonstrate the reliable performance of our algorithm in high dimensions.

▷ *CLARSTA: A random subspace trust-region algorithm for convex-constrained derivative-free optimization* [11].

Non-Convergence Analysis: What, Why, and How?

Zaikun Zhang *Sun Yat-sen University*
with Cunxin Huang (The Hong Kong Polytechnic University)

15:45–16:10 ↑ Back to schedule

We present a non-convergence theory for probabilistic direct search, a randomized derivative-free optimization method, where non-convergence means the failure to produce iterates that achieve stationarity asymptotically. The motivation is to understand whether the submartingale-like assumption in the existing convergence theory is essential or merely an artifact of the analysis techniques. For convex objectives, we prove that the probability of non-convergence is positive, provided that the polling directions satisfy a probabilistic ascent condition that is essentially the opposite of the submartingale-like convergence condition. Furthermore, we establish a lower bound for the non-convergence probability. For the typical implementation of this method, where each iteration draws a fixed number of random polling directions independently and uniformly from the unit sphere, our theory implies that the method is not globally convergent if the number of directions is below the threshold specified in the convergence theory, and the submartingale-like assumption is confirmed to be essential for convergence. Our theory is obtained by examining two random series that control the distance from any iterate to the starting point and estimating the probability for these series to stay below certain bounds.

▷ *Non-convergence Analysis of Probabilistic Direct Search* [26].

Birkhoff interpolation models for optimization with some available derivatives

Jeffrey Larson *Argonne National Laboratory*
with Matt Menickelly and Evan Toler (Argonne National Laboratory)

16:20–16:45 ↑ Back to schedule

Many modern optimization problems sit in an awkward middle ground: full gradients are unavailable or too costly, yet some derivative information is accessible for select variables. This talk presents a trust-region framework to address such “some available derivatives” settings. Our approach builds local models using Birkhoff interpolation, combining function values with the available partial derivatives. We introduce a practical notion of poisedness for such mixed data sets and describe procedures to maintain well-conditioned sets for interpolation. Our approach aims to reduce expensive objective queries, improve robustness relative to value-only model building, and remain reliable when derivative availability varies across iterations or coordinates.

▷ *Birkhoff interpolation models for optimization with some available derivatives* [30].

Hybrid Neural-Quadratic Model-Based Derivative-Free Optimization

Stefan Wild *Lawrence Berkeley National Lab*
with Pengcheng Xie (Lawrence Berkeley National Laboratory)

16:55–17:20 ↑ Back to schedule

Quadratic models are popular choices in model-based DFO because they can capture local nonlinearities while remaining amenable to efficient trust-region subproblem solution. Neural network models, by contrast, excel at capturing more global, higher-order nonlinearity. In this talk we attempt to get the best of both worlds through a new hybrid surrogate framework for expensive DFO that integrates neural network regression and classical quadratic models within a unified trust-region methodology. Key to our framework is adaptive model selection. Numerical results on benchmark test sets show that model switching is essential — different surrogates dominate at different stages — and that their combination yields consistently better performance than any fixed surrogate choice.

Tuesday, June 23 — Morning

Adaptive Optimization Methods and Noisy Oracles

Albert Berahas *University of Michigan*
with Katya Scheinberg (Georgia Tech) and Liyuan Cao (Nanjing University)

9:00–9:25 ↑ Back to schedule

In this talk, we discuss adaptive optimization methods, including line-search and trust-region frameworks, in the presence of deterministic, stochastic, and probabilistic noise. We present unified algorithmic frameworks, convergence guarantees, and complexity results, along with illustrative numerical experiments. This work is based on a forthcoming review paper on noisy oracle optimization in collaboration with Professors Cao and Scheinberg.

Chance-constrained blackbox optimization

Tudor Buhut *GERAD, Polytechnique Montréal*
with Antoine Lesage-Landry, Sébastien Le Digabel (GERAD, Polytechnique Montréal)

9:35–10:00 ↑ Back to schedule

Stochastic blackbox optimization (BBO) is a field that sees discord between the theoretical analysis of its algorithms and their actual implementations. Numerous algorithms have been shown to almost surely converge under assumptions that require sampling the stochastic blackbox an excessive number of times as the step size of the algorithm decreases. This is infeasible in practice and is therefore never done. This work introduces a stochastic BBO framework that embeds practical considerations within its theoretical analysis. We present a solver-agnostic wrapper that minimizes the median of a stochastic objective function subject to chance constraints. Our wrapper is compatible with Bayesian optimization, direct search methods and most other non-stochastic BBO solvers. Confidence intervals establish the probability that the probability of satisfaction is greater than a user-defined threshold as well as the probability that the median of the objective sees descent from the incumbent to the candidate point provided by the backbone solver. For any fixed confidence and at any given iteration of the blackbox solver, the number of samples to be queried is bounded. The probability that the wrapper decides correctly at each iteration is thus known, all while avoiding prohibitive sampling. This fact is employed to discuss convergence results inherited from the backbone solver. Computational results using NOMAD as the solver illustrate the performance and practicability of our approach.

Errors in stochastic derivative free optimization

Katya Scheinberg *Georgia Institute of Technology*
with Albert Berahas (University of Michigan) and Liyuan Cao (Nanjing University)

10:10–10:35 ↑ Back to schedule

We will discuss how different types of errors arise in stochastic derivative free optimization when interpolation or finite different techniques are combined with random subspaces and stochastic function evaluations with and without common random numbers. We will classify the different types of errors and present their effect on the algorithmic behaviour.

An Extended Kolmogorov Inequality for Sharper Complexity Bounds in Randomized Block Coordinate Descent

Zao Jiang *School of Mathematics, Sun Yat-sen University*
with Cunxin Huang (The Hong Kong Polytechnic University), Zaikun Zhang (Sun Yat-sen University)

11:05–11:30 ↑ Back to schedule

We derive a sharper complexity bound for randomized block coordinate descent by generalizing Kolmogorov's inequality. The inequality is extended to random variables satisfying a weak conditional version of Cramer's condition, removing the need for independence and recovering the classical bound as a special case. This condition is satisfied by bounded random variables and, more generally, by those with finite second moments, and yields a one-sided tail bound. For smooth functions with block-wise Lipschitz gradients, our bound eliminates the multiplicative dependence between confidence and accuracy, scaling only with their maximum. The extended inequality is not limited to randomized block coordinate descent. It provides a general tool for high-probability analysis of randomized iterative methods.

A Riemannian framework for feasible DFO with equality constraints

Sébastien Le Digabel *GERAD, Polytechnique Montréal*
with Charles Audet and Sacha Benarroch-Lelong (GERAD, Polytechnique Montréal)

11:40–12:05 ↑ Back to schedule

Solving derivative-free optimization problems with equality constraints is especially challenging when nonfeasible iterates may cause objective function evaluations to fail. When these constraints are a priori, knowledge of the geometry of the feasible set can help guide the optimization in a parametrization-like approach. In the case when this set is a Riemannian manifold, subproblems can be formulated in its tangent spaces, and the iterates brought back to the manifold by a so-called retraction. This approach implicitly enforces dimensionality reduction while yielding only feasible iterates. While this fact is known from the literature, the main contribution of this work is the introduction of a formal Riemannian framework compatible with any derivative-free solver, that automates computations of tangent spaces and solving the subproblems. When used with a Mesh Adaptive Direct Search solver, we prove that this framework yields theoretical guarantees in terms of Clarke's generalized derivatives in the tangent spaces.

Tuesday, June 23 — Afternoon

Complexity Analysis of Regular Simplicial Search for Minimizing Smooth Black-Box Functions

Liyuan Cao *Nanjing University*
with Wei Hu (AMSS, Chinese Academy of Sciences) and Jinxin Wang (Department of Statistics, University of Chicago)

13:30–13:55 ↑ Back to schedule

Simplex-type methods, such as the classical methods of Spendley et al. and Nelder-Mead, are among the most widely used derivative-free optimization algorithms in practice. Their theoretical complexity, however, remains mysterious. In this project, we consider a simplified algorithm that starts with a regular simplex and only uses the reflection and the shrink operations. Under the assumption that the objective function is smooth with Lipschitz continuous gradient, we establish the query complexity for this algorithm to find an epsilon-stationary point. We also derive its query complexity for reducing the optimality gap below epsilon when the objective is convex, and for reducing the distance to the optimal solution below epsilon when the objective is strongly convex.

- ▷ *The error in multivariate linear extrapolation with applications to derivative-free optimization* [9].

Zeroth-Order Min-Max Optimisation for Nonconvex-Nonconcave Functions

Youssef Diouane *GERAD, Polytechnique Montréal*
with Amir Ali Farzin, Yuen Man Pun, Philipp Braun (Australian National University), Antoine Lesage-Landry (GERAD, Mila, Polytechnique Montréal), Iman Shames (The University of Melbourne)

14:05–14:30 ↑ Back to schedule

In this talk, we investigate the performance of the random Gaussian smoothing Zeroth-Order ExtraGradient (ZO-EG) method for min-max optimization problems with possibly nonconvex-nonconcave (NC-NC) objective functions. We consider unconstrained and constrained settings, as well as differentiable and nondifferentiable cases. We study these min-max problems from the perspective of variational inequalities. For the unconstrained case, we establish the convergence of the ZO-EG algorithm to a neighborhood of a stationary point of the NC-NC objective function, show how the radius of this neighborhood can be controlled through variance reduction, and derive the corresponding complexity results. Numerical comparisons on several machine learning tasks against existing direct-search methods for min-max problems highlight the potential of the proposed approach, particularly in the later stages of optimization.

▷ *Min-Max Optimisation for Nonconvex-Nonconcave Functions Using a Random Zeroth-Order Extragradient Algorithm* [20].

Solving the Offline and Online Min-Max Problem of Non-smooth Submodular-Concave Functions: A Zeroth-Order Approach

Iman Shames *The University of Melbourne*
with Amir Ali Farzin, Yuen Man Pun, Philipp Braun (Australian National University), Tyler Summers (University of Texas at Dallas)

14:40–15:05 ↑ Back to schedule

In this talk, I'll present recent results on optimizing max–min and min–max problems where the objective is concave in the maximizer but submodular–and possibly non-smooth–with respect to the minimizer. These problems appear in various learning and decision-making applications, yet classical smooth optimization tools fail to apply directly. I'll describe a zeroth-order method that combines two key ingredients: the subgradient of the Lovász extension to handle submodularity on the minimization side, and Gaussian smoothing to estimate gradients for the maximization variable. We'll see that, in the offline setting, the algorithm converges in expectation to an ϵ-saddle point, and in the online setting, it achieves an $O(\sqrt{N\bar{P}_N})$ bound on the expected duality gap, where $\bar{P}_N$ measures the variability of the optimal responses. I'll also discuss the complexity guarantees, how to tune hyperparameters in practice, and show some experimental results that illustrate the theoretical findings.

Multi-Fidelity Warm-Start Optimization: Theory and Application

Zhongda Huang *UBC Okanagan*
with Warren Hare (UBC Okanagan)

15:45–16:10 ↑ Back to schedule

Optical computed tomography (CT) scanners are powerful imaging tools used in three-dimensional radiation dosimetry. Solid-tank scanners use an acrylic block to simplify setup. We examine a blackbox multi-fidelity multiobjective problem coming from solid tank design. We study methods to exploit the multi-fidelity nature of the problem to enhance algorithm speed and solution quality. Motivated by this problem, we develop a theoretical framework for multi-fidelity warm-start optimization under smooth and strongly convex objectives. we establish the optimal allocation of iterations between fidelity levels, including closed-form cost-minimizing splits.

▷ *A multiobjective multifidelity optimization approach for the design of scanners for use in optical computed tomography imaging* [27].

Derivative-Free Inexact Mirror Langevin Algorithms

Raghu Pasupathy *Purdue University, Department of Statistics*
with Antik Chakraborty (Purdue University, Department of Statistics)

16:20–16:45 ↑ Back to schedule

We develop a derivative-free mirror Langevin framework for sampling and optimization over constrained convex domains equipped with barrier geometry. Mirror Langevin methods combine mirror descent and Langevin dynamics to enable efficient sampling and optimization, but existing theory assumes direct access to gradients. In many important applications including simulation optimization, stochastic network control, and black-box experimental design, gradients exist but are not observable and must be estimated from noisy function evaluations. Building on classical zeroth-order techniques such as finite differences, SPSA, and Gaussian smoothing, we introduce derivative-free mirror Langevin algorithms that integrate randomized gradient estimation with barrier-based proximal dynamics. Our analysis shows that the self-concordant geometry of barrier functions naturally controls gradient estimation error and inexact proximal updates, yielding explicit recursive bounds and convergence guarantees in Wasserstein distance. These results provide a unified foundation for derivative-free mirror Langevin methods and enable scalable algorithms for constrained stochastic optimization and sampling in high-dimensional, simulation-driven settings.

Function-free Optimization via Comparison Oracles

Zikai Xiong *Georgia Tech*
with Katya Scheinberg (Georgia Institute of Technology)

16:55–17:20 ↑ Back to schedule

In this work, we study optimization specified only through a comparison oracle: given two points, it returns which one is preferred. We call it function-free optimization because we do not assume access to, nor the existence of, a canonical objective function; instead, our goal is to find points that are not preferred over by any other point, which we denote as the optimal solutions. This model arises in preference- and ranking-based settings, where the objective values and derivatives are unavailable, meaningless, or non-identifiable. We develop a systematic, analytical, and algorithmic framework based on the geometry of the preference level set, as it remains well-defined from comparisons alone. We introduce a new optimality measure, called the "level set optimality gap," which denotes the distance between the level set and the optimal solutions. We then propose a method for estimating normal directions to level sets with complexity $O(d \cdot \log(1/\epsilon))$, and show that the method is nearly optimal. Leveraging these normal direction estimates, we show that a normal direction descent method reaches level-set optimality gap at most ϵ in $O(d/\epsilon^2)$ iterations when the level sets are smooth and convex. The underlying representative function of our algorithm may be nonsmooth, nonconvex, or even discontinuous. Since prior knowledge in practical applications of function-free optimization is usually very limited, we also give parameter-free schemes for normal direction estimation and solving the optimization problem with nearly the same order of complexity.

▷ *Function-free Optimization via Comparison Oracles* [41].

Thursday, June 25 — Morning

Polling Set Construction and Worst-Case Complexity for Direct Search under Polyhedral Convex Constraints

Lindon Roberts *University of Melbourne*
with Clément Royer (Université Paris Dauphine-PSL)

9:00–9:25 ↑ Back to schedule

Direct search is one of the most popular derivative-free optimization (DFO) paradigms [29]. To analyze and implement direct search, one typically relies on positive spanning sets, a quite separate notion from ideas used in the analysis of interpolation sets in model-based DFO. In this talk, we will introduce a new theoretical underpinning for direct-search methods, that is coherent with model-based DFO theory and can be defined for general convex constraints. For polyhedral convex constraints, we are able to construct polling sets that meet our new theoretical requirements and gives some theoretical justification to existing practical heuristics [22, 31]. Our numerical results confirm that adding these extra directions significantly improves practical performance over both existing theoretical poll set constructions and existing heuristics. This talk is based on [39].

- ▷ *Direct search based on probabilistic feasible descent for bound and linearly constrained problems* [22].
- ▷ *Optimization by direct search: New perspectives on some classical and modern methods* [29].
- ▷ *Implementing Generating Set Search Methods for Linearly Constrained Minimization* [31].
- ▷ *Polling Set Construction and Worst-Case Complexity for Direct Search under Polyhedral Convex Constraints* [39].

Practical computation of cosine measures

Warren Hare *University of British Columbia*
with Scholar Sun (Georgia Tech)

9:35–10:00 ↑ Back to schedule

Since you are at DFOS, you might already know what cosine measures are. If I'm talking later in the week, then you almost surely know what cosine measures are. But, do you know how to compute them? In R^{100}? In less than a minute? (P.S. If no one else tells you first, then I promise I'll tell you what cosine measures are.)

It's a graph! It's a polytope! No, it's a PSS!

Clément Royer *Université Paris Dauphine-PSL*
with Denis Cornaz and Sébastien Kerleau (Université Paris Dauphine-PSL)

10:10–10:35 ↑ Back to schedule

Positive spanning sets (PSSs) have played a major role in designing derivative-free optimization methods, particularly of direct-search type. Most implementations rely on a handful of PSS examples, that are almost exclusively positive bases, i.e. inclusionwise minimal PSSs. While those specific examples are well understood by the community, a more general understanding of positive bases (and PSSs) remains an active topic of research. In this talk, I will highlight connections between positive spanning sets, graphs, and polytopes. I will first explain how positive spanning sets relate to strongly connected digraphs, and illustrate how the ear decomposition from graph theory allows to characterize positive bases. I will then turn to polytope theory, that provides a perhaps unexpected way of constructing PSSs and positive bases of smallest cardinality. Finally, I will hint at several open questions around PSSs, where graphs and polytopes might offer a novel perspective.

- ▷ *A characterization of positive spanning sets with ties to strongly connected digraphs* [15].
- ▷ *Using orthogonally structured positive bases for constructing positive k-spanning sets with cosine measure guarantees* [24].
- ▷ *Graphical and geometrical perspectives on positive spanning sets with applications to derivative-free optimization* [28].

Convergence towards a local solution by algorithms with a covering step

Pierre-Yves Bouchet *Polytechnique Montréal*
with Charles Audet (GERAD, Polytechnique Montréal) and Loïc Bourdin (XLIM Research Institute)

11:05–11:30 ↑ Back to schedule

This talk introduces a new step to be added to most optimization algorithms to strengthen their convergence analyses. By design, this so-called covering step may ensure that all accumulation points of the sequence of incumbent solutions generated by the algorithm are local solutions to the optimization problem, within a ball of pre-defined radius that is chosen as an algorithmic parameter. This new result holds true even for a discontinuous objective function, under a mild assumption that we discuss in details. We also provide a practical construction scheme for the covering step that works at low additional cost per iteration, and we illustrate in details how to fit this covering step on the class of Direct Search Method algorithms. Numerical experiments are also provided to highlight the contexts where the covering step yields practical improvements in addition to theoretical ones.

- ▷ *Convergence towards a local minimum by direct search methods with a covering step* [1].
- ▷ *A partitioned optimization framework for structure-aware problems* [2]
- ▷ *On the computation of the cosine measure in high dimensions* [25].

Black-Box Optimization with Hexaly

Fred Gardi *Hexaly*

11:40–12:05 ↑ Back to schedule

In black-box optimization problems, the analytical form of the objective function and constraints is unknown, and evaluating new solutions can be costly in terms of time and memory. Hexaly provides a surrogate modeling method for optimizing these problems using radial basis function (RBF) models, which helps to reduce the number of evaluations by carefully selecting the points to evaluate. The algorithm combines exploration and intensification phases to make the most of a limited evaluation budget. Additionally, Hexaly handles both analytical and black-box constraints commonly found in industrial applications and can optimize multiple black-box objectives simultaneously. This presentation presents the approach implemented in Hexaly and illustrates its effectiveness in real-world applications.

Thursday, June 25 — Afternoon

Margherita Porcelli *University of Florence*
with Filippo Marini (Università degli Studi di Firenze) and Elisa Riccietti (ENS de Lyon)

13:30–13:55 ↑ Back to schedule

We present a new multilevel stochastic framework for the solution of optimization problems. The proposed approach uses random regularized first-order models that exploit an available hierarchical description of the problem, being either in the classical variable space or in the function space, meaning that different levels of accuracy for the objective function are available. We present the converge analysis of the method and show its numerical behavior on the solution of finite-sum minimization problems arising in binary classification problems.

- ▷ *A multilevel stochastic regularized first-order method with application to finite sum minimization, Mathematics of Computation* [34].

Blackbox multi-period power flow optimization under realistic operational constraints

Christophe Tribes *GERAD, Polytechnique Montréal*
with Martin De Montigny, Ludovic Salomon, and Hussein Suprême (Hydro-Québec)

14:05–14:30 ↑ Back to schedule

This work presents a scalable blackbox optimization (BBO) approach for multi-period power flow planning under realistic operational constraints. An AC power flow simulation captures detailed temporal system behavior in a quasi-static time-series environments. The overall problem is high-dimensional with mixed discrete–continuous variables and contains numerous operational constraints. A novel decomposition strategy splits the multi-period problem into loosely coupled single-period BBO subproblems solved in parallel, with iterative coordination to enforce temporal consistency. We leverage a derivative-free algorithm (MADS with progressive barrier) to handle relaxable constraints and identify feasible or infeasible yet promising solutions. The BBO subproblems integrate soft constraints on voltages and switching actions to reflect practical operational guidelines. Experiments on a large-scale transmission system covering seven days of operation demonstrate substantial loss reductions while preserving voltage security and limiting switching activity, illustrating an effective balance between performance and operational feasibility.

A Direct Multisearch Approach (DMS) for Many-Objective Derivative-Free Optimization

Everton José da Silva *NOVA School of Science and Technology*
with Ana Luisa Custodio (Universidade Nova de Lisboa)

14:40–15:05 ↑ Back to schedule

DMS was originally developed for continuous constrained multiobjective derivative–free optimization, with a general number of objective function components. Still, its performance was never tested for problems with more than three objectives. We propose DMS-Reduction, a variant that uses reduction methods, such as sketching techniques and correlation approaches, to reduce the number of objective function components and related variables to be addressed at each iteration. We detail the algorithmic structure and report promising numerical results in tackling many-objective optimization problems.

- ▷ *Direct multisearch for multiobjective optimization* [16].
- ▷ *An inexact restoration direct multisearch filter approach to multiobjective constrained derivative-free optimization* [42].

Data-driven optimization: efficient adaptive learning for self-driving laboratories

Nick Sahinidis *Georgia Tech*

15:45–16:10 ↑ Back to schedule

Self-driving laboratories promise to compress materials-discovery timelines from years to weeks by replacing trial-and-error experimentation with closed-loop, algorithm-guided campaigns. Yet, despite the rapid proliferation of robotic and automation hardware, today's autonomous labs rely almost exclusively on Bayesian optimization (BO) to decide what experiment to run next. BO is a sensible approach to low-dimensional optimization problems with smooth response surfaces, but it struggles in precisely the regimes that matter most for real materials campaigns: tight experimental budgets, dozens of process parameters, mixed-integer choices, hard physical constraints, and noisy, expensive measurements.

In this talk, I will show how moving from BO to partitioning-based algorithms can substantially improve data efficiency, scale gracefully to dozens of process variables, and handle the constraints and noise that characterize realistic experimental campaigns. I will summarize a recently completed large-scale black-box optimization (BBO) benchmark in which we compared 42 solvers across 502 problems ranging from one to 300 dimensions and from smooth and convex to nonsmooth and nonconvex. The results overturn several community assumptions: BO solves only about 9% of problems within a 2,500-evaluation budget, while a new branch-and-model (BAM) algorithm reaches an 81% success rate, with GLCCLUSTER, MULTIMIN, MCS, and SNOBFIT also performing strongly. A minimal, irreducible set of eight complementary solvers attains 88% solvability on the full suite.

I will then move from in-silico benchmarks to the wet lab, presenting a recent algorithm-guided experimental campaign on high-performance perovskite solar cells in which a non-BO solver was used to co-optimize six process variables spanning the perovskite, electron-transport, and hole-transport layers. Time permitting, I will also share early results from applying ensembles of BBO algorithms to digital twins of self-driving labs across additional materials systems.

I will close with a forward-looking research vision: accelerating autonomous labs by developing, benchmarking, and experimentally validating data-efficient adaptive algorithms across batteries, semiconductors, catalysts, polymeric membranes, and biomolecules. The benchmarking software will be released as open source, with BAM and most BBO software available free to academic users, so that experimental groups can deploy these tools on their own self-driving platforms.

▷ *A large-scale benchmarking of deterministic and stochastic derivative-free optimization algorithms* [35].

Black-box Active Learning for Discovering New Explanations in Complex System Exploration

Paul Saves *Université de Toulouse Capitole & IRIT*
with Nicolas Verstaevel, Benoit Gaudou (IRIT & Université Toulouse Capitole)

16:20–16:45 ↑ Back to schedule

Exploring the parameter space of socio-technical complex systems, such as Agent-Based Models, presents significant challenges due to high dimensionality and computational cost. As these simulators typically function as non-differentiable black-box models, effective exploration requires derivative-free strategies. However, traditional sampling methods, such as Stepwise Uncertainty Reduction or Static Design of Experiments, typically prioritize minimizing global prediction error or output variance. These approaches often fail to capture the diverse causal mechanisms underlying system behavior, missing critical instances where different input combinations yield similar outputs for fundamentally different reasons.

In this work, we propose a novel Active Learning framework designed to discover "new explanations" rather than simply accumulating data points. Our model-agnostic method integrates Surrogate Modeling with post-hoc explainable artificial intelligence to guide the exploration process. The methodology proceeds in an iterative Bayesian loop, maximizing the explanatory diversity within the collected data. By explicitly targeting points that are distinct in their explanatory structure (feature contribution profiles), our framework efficiently uncovers rare behavioral regimes and distinct internal logic. Consequently, this framework achieves a deep structural understanding with fewer simulations than standard prediction-based strategies. It proves particularly effective in characterizing critical system behaviors, including tipping points.

▷ *Multi-Fidelity Constrained Bayesian Optimization with Application to Aircraft Wing Design* [14].

▷ *A Stepwise uncertainty reduction approach to constrained global optimization* [37].

▷ *Early-warning signals for critical transitions* [40].

Multi-objective Bayesian optimization meets multiple gradient descent

Mickael Binois *Inria*
with Jean-Antoine Désidéri (Inria)

16:55–17:20 ↑ Back to schedule

Designed for limited evaluation budgets, multi-objective Bayesian optimization (BO) typically targets finding the entire Pareto front. With more than a few objectives, it is more realistic to target specific parts of the Pareto front. One such way is to rely on the multiple gradient descent principle, improving on all objectives at once from a given starting point. For moderately large input dimensions, the combination of BO with trust-region proved to be efficient, but mostly in the mono-objective context. Following this setup, we propose a trust-region based multi-objective Bayesian optimization framework relying on the multiple gradient descent principle. We show the benefit of the approach on several examples.

Friday, June 26 — Morning

Surrogate-based categorical neighborhoods for mixed-variable blackbox optimization

Edward Hallé-Hannan *GERAD, Polytechnique Montréal*
with Charles Audet, Youssef Diouane, Sébastien Le Digabel, Christophe Tribes (GERAD, Polytechnique Montréal)

9:00–9:25 ↑ Back to schedule

In simulation-based engineering, design choices are often obtained through the optimization of complex blackbox models. These problems frequently involve mixed-variable domains with quantitative and categorical variables. Unlike quantitative variables, categorical variables lack an inherent structure, which makes them difficult to handle, especially in the presence of constraints. This work proposes a systematic approach to structure and model categorical variables in constrained mixed-variable blackbox optimization. Surrogate models of the objective and constraint functions are used to induce problem-specific categorical distances. From these distances, surrogate-based neighborhoods are constructed using notions of dominance from bi-objective optimization, so that information from both the objective and the constraint functions is jointly considered. As a proof of concept, these neighborhoods are employed within CatMADS, where the problem-specific distances are derived from Gaussian processes. The resulting method is benchmarked on the Cat-Suite collection of 60 mixed-variable optimization problems and compared against state-of-the-art solvers. Data profiles indicate competitive performance of the proposed method. Finally, its performance is demonstrated on a real-world aircraft design application.

- ▷ *CatMADS: Mesh Adaptive Direct Search for constrained blackbox optimization with categorical variables* [5].
- ▷ *A distance for mixed-variable and hierarchical domains with meta variables* [23].

Boosting BFO: a new strategy to efficiently explore the integer lattice

Federico Grillini *University of Modena and Reggio Emilia*
with Margherita Porcelli (University of Florence), Francesco Rinaldi (University of Padova)

9:35–10:00 ↑ Back to schedule

A common difficulty in mixed-integer blackbox optimization is to balance the exploration of discrete and continuous variables. In this talk, we propose a trust region framework designed by combining the algorithmic structure of the Brute Force Optimizer (BFO) with the use of primitive directions. Furthermore, we employ information from past function evaluations to bias the search toward promising regions of the discrete domain. Finally, the numerical experiments show the reliability and efficiency of the proposed method, when compared to state of the art solvers.

- ▷ *BFO, A Trainable Derivative-free Brute Force Optimizer for Nonlinear Bound-constrained Optimization and Equilibrium Computations with Continuous and Discrete Variables* [38].

Adaptive direct search algorithms with relaxable and quantifiable constraints

Théo Denorme *GERAD, Polytechnique Montréal*
with Charles Audet, Youssef Diouane, Sébastien Le Digabel, Christophe Tribes (GERAD, Polytechnique Montréal)

10:10–10:35 ↑ Back to schedule

This talk introduces ADS-PB, an extension of the Adaptive Direct Search (ADS) framework to constrained blackbox optimization problems with quantifiable and relaxable constraints. The problems considered involve minimizing an objective function subject to inequality constraints, where both the objective and constraint functions are available only through blackbox evaluations. In contrast with the extreme barrier approach used in ADS, the proposed framework incorporates a progressive barrier mechanism, allowing infeasible points to be evaluated and exploited through a measure of constraint violation. ADS-PB is thus suited to problems where constraint violations can be quantified and progressively reduced. The resulting algorithm does not rely on mesh structures or sufficient decrease conditions, and a convergence analysis is provided under standard assumptions of derivative-free optimization. Its performance is assessed on constrained test problems and compared with existing methods. The integration within the NOMAD software is also discussed.

▷ *Adaptive direct search algorithms for constrained optimization* [3].

Managing Multiple Models in Expensive Derivative-Free Optimization with Contextual Bandits

Matt Menickelly *Argonne National Laboratory*

11:05–11:30 ↑ Back to schedule

In computational science workflows, it is often the case that 1) simulations do not permit readily the algorithmic computation of derivatives, 2) objective functions for optimization involve multiple simulation outputs, and 3) those simulation outputs can be performed (at least partially) in parallel. While the Argonne-maintained derivative-free optimization (DFO) solvers in the Interpolation-Based Composite Derivative-Free Optimization (IBCDFO) repository are specifically designed to address problems exhibiting the first two of these characteristics, they do not support the third at the time of writing this abstract. In this work, we reexamine our past work on a class of randomized algorithms, stochastic average model (SAM) methods. SAM methods are conceptually similar to stochastic average gradient methods, and effectively require that only randomized subsets of simulation outputs be locally modeled in each iteration of a model-based DFO method. From the perspective of variance reduction, given a batch size, the optimal sampling distribution for a batch within a SAM method permits a closed-form, but the closed-form distribution contains unknowable parameters. However, clues concerning the values of the unknowable parameters are often available either through domain expertise, alternative lower-fidelity simulations, or simply regression models. Motivated by this, we consider the problem of determining sampling distributions for SAM methods as a contextual bandit problem and we apply the Exponential weights algorithm for Exploration and Exploitation with Experts (Exp4). Results are demonstrated on both synthetic problems and a problem of density functional theory calibration.

On Complexity of Model-Based Derivative-Free Methods

Abraar Chaudhry *Georgia Tech*
with Katya Scheinberg (Georgia Tech)

11:40–12:05 ↑ Back to schedule

In many applications of mathematical optimization, one may wish to optimize an objective function without access to its derivatives. These situations call for derivative-free optimization (DFO) methods. Among the most successful approaches in practice are model-based trust-region methods, such as those pioneered by M.J.D Powell. While relatively complex to implement, these methods are now available in standard scientific computing platforms, including MATLAB and SciPy. However, theoretical analysis of their computational complexity lags behind practice. In particular, it is important to bound the number of function evaluations required to achieve a desired level of accuracy. In this paper we systematically derive complexity bounds for classical model-based trust-region methods and their modern variations. We establish, for the first time, that these methods can have the same worst case complexity than any other known DFO method.

▷ *On Complexity of Model-Based Derivative-Free Methods* [10].

Friday, June 26 — Afternoon

Extrapolation-based Direct Search for Nonsmooth Stochastic Zeroth-Order Optimization

Anthony Palmieri *University of Padova*
with Francesco Rinaldi (University of Padova), Sara Shashaani (North Carolina State University)

13:30–13:55 ↑ Back to schedule

We propose and analyze a direct-search scheme for unconstrained stochastic zeroth-order minimization of locally Lipschitz, possibly nonsmooth, objectives that combines a stochastic line search with random dense search directions. We impose minimal assumptions on the stochastic oracle, that is, zero-mean and finite-variance noise with no variance-stepsize coupling. Under these conditions, we prove almost-sure convergence to Clarke stationary points, while also establishing some preliminary complexity results. Numerical experiments on a DFO benchmark suite highlight competitive performance against well-established stochastic direct-search methods.

▷ *Extrapolation-based Direct Search for Nonsmooth Stochastic Zeroth-Order Optimization* [36].

Benchmarking Techniques for Bilevel Derivative-free Optimization

Valentin Dijon *GERAD, Polytechnique Montréal*
with Youssef Diouane and Charles Audet (GERAD, Polytechnique Montréal)

14:05–14:30 ↑ Back to schedule

This talk introduces a rigorous methodology to compare derivative-free optimization (DFO) algorithms for bilevel optimization problems.

Issues related to the admissibility of solutions or the effort spent to solve the lower-level problem are not always considered in the literature. Updated DFO benchmarking techniques are proposed to account for the specific challenges arising from bilevel optimization. First, the computational cost devoted to solving the lower-level problem is explicitly incorporated, allowing for a more faithful representation of the overall effort required by each solver to compute a solution. Then, a refereeing process is introduced to assess the admissibility of candidate solutions, thereby ensuring they lie in the induced region.

Numerical experiments are conducted on problems from the BOLIB test set using the Julia programming language. The resulting profiles and plots highlight the impact on bilevel DFO algorithms performance when second-level effort and optimality are included in benchmarking techniques.

▷ *Benchmarking Bilevel Derivative-Free Optimization Algorithms* [4].

Benchmarking algorithms on large test sets

Charles Audet *GERAD, Polytechnique Montréal*
with Christophe Tribes (GERAD, Polytechnique Montréal) and Warren Hare (University of British Columbia)

14:40–15:05 ↑ Back to schedule

Benchmarking is essential for assessing the effectiveness of optimization algorithms. This is especially true in derivative-free optimization, where target problems are often complex simulations that require extensive time to evaluate. This limits the number of evaluations that can be performed, making it critical to have a good understanding of the potential quality of various algorithms. This talk reviews standard benchmarking methods, including convergence plots, performance profiles, data profiles, and accuracy profiles, widely used to evaluate optimization algorithms. The primary contribution is a formal extension of these benchmarking techniques to three specific contexts: constrained optimization, multi-objective optimization, and surrogate-based optimization. Benchmarking codes are proposed and made available to the community.

▷ *Derivative-Free and Blackbox Optimization (2nd edition)* [6].

▷ *A summary of benchmarking constrained, multi-objective and surrogate-assisted optimization methods* [7].

References

[1] C. Audet, P.-Y. Bouchet, and L. Bourdin. Convergence towards a local minimum by direct search methods with a covering step. *Optimization Letters*, 19(1):211–231, 2025. doi:10.1007/s11590-024-02165-2.

[2] C. Audet, P.-Y. Bouchet, and L. Bourdin. A partitioned optimization framework for structure-aware problems. *Journal of Optimization Theory and Applications*, 210(1):9, 2026. doi:10.1007/s10957-026-03042-x.

[3] C. Audet, T. Denorme, Y. Diouane, S. Le Digabel, and C. Tribes. Adaptive direct search algorithms for constrained optimization. Technical Report G-2025-53, Les cahiers du GERAD, 2025. doi:10.48550/arXiv.2507.23054.

[4] C. Audet, V. Dijon, and Y. Diouane. Benchmarking Bilevel Derivative-Free Optimization Algorithms. Technical Report G-2026-26, Les cahiers du GERAD, 2026. doi:10.48550/arXiv.2605.30531.

[5] C. Audet, Y. Diouane, E. Hallé-Hannan, S. Le Digabel, and C. Tribes. CatMADS: Mesh Adaptive Direct Search for constrained blackbox optimization with categorical variables. Technical Report G-2025-42, Les cahiers du GERAD, 2025. doi:10.48550/arXiv.2506.06937.

[6] C. Audet and W. Hare. *Derivative-Free and Blackbox Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, Cham, Switzerland, 2nd edition, 2026. doi:10.1007/978-3-032-00906-7.

[7] C. Audet, W. Hare, and C. Tribes. A summary of benchmarking constrained, multi-objective and surrogate-assisted optimization methods. *Optimization Letters*, 2026. doi:10.1007/s11590-026-02302-z.

[8] M. Bergamaschi, S. Venturini, F. Tudisco, and F. Rinaldi. An efficient network-aware direct search method for influence maximization. *Journal of Complex Networks*, 13(6):cnaf042, 2025. doi:10.1093/comnet/cnaf042.

[9] L. Cao, Z. Wen, and Y.X. Yuan. The error in multivariate linear extrapolation with applications to derivative-free optimization. *Optimization Methods and Software*, pages 1–45, 2026. doi:10.1080/10556788.2026.2678174.

[10] A. Chaudhry and K. Scheinberg. On Complexity of Model-Based Derivative-Free Methods. Technical Report 2510.14935, arXiv, 2025. doi:10.48550/arXiv.2510.14935.

[11] Y. Chen, W. Hare, and A. Wiebe. CLARSTA: A random subspace trust-region algorithm for convex-constrained derivative-free optimization. Technical Report 2506.20335, arXiv, 2026. doi:10.48550/arXiv.2506.20335.

[12] R. Filomeno Coelho, C. Sainvitu, and T. Benamara. Dimensionality reduction combining variable grouping and UMAP for structural optimization. *Structural and Multidisciplinary Optimization*, 68(5):106, 2025. doi:10.1007/s00158-025-04028-5.

[13] A.R. Conn, K. Scheinberg, and L.N. Vicente. *Introduction to Derivative-Free Optimization*. MOS-SIAM Series on Optimization. SIAM, Philadelphia, 2009. doi:10.1137/1.9780898718768.

[14] O. Cordelier, Y. Diouane, N. Bartoli, and E. Laurendeau. Multi-Fidelity Constrained Bayesian Optimization with Application to Aircraft Wing Design, 2025. doi:10.2514/6.2025-3474.

[15] D. Cornaz, S. Kerleau, and C.W. Royer. A characterization of positive spanning sets with ties to strongly edge-connected digraphs. *Discrete Applied Mathematics*, 374:105–119, 2025. doi:10.1016/j.dam.2025.05.025.

[16] A.L. Custódio, J.F.A. Madeira, A.I.F. Vaz, and L.N. Vicente. Direct multisearch for multiobjective optimization. *SIAM Journal on Optimization*, 21(3):1109–1140, 2011. doi:10.1137/10079731X.

[17] M. Diago, X. Lebeuf, T.J.M. Meango, A. Côté, and C. Rajotte. Risk-Based Optimization of Periodic Maintenance for Power Grid Equipment. In *2025 Annual Reliability and Maintainability Symposium (RAMS)*, pages 1–6. IEEE, 2025. doi:10.1109/RAMS48127.2025.10935127.

[18] Q. Dollon, E. Ghorbani, and F.P. Gosselin. Bayesian input-state-parameter inference of hydrodynamic bearings: From partial displacement measurements to force reconstruction. *Mechanical Systems and Signal Processing*, 230:112610, 2025. doi:10.1016/j.ymssp.2025.112610.

[19] M. Fahed, A. Favrel, M. Gagnon, S. Alarie, and M. Kokkolaras. Hydropower unit digital twin calibration using monitoring data and blackbox optimization. *IOP Conference Series: Earth and Environmental Science*, 1561(1):012018, 2025. doi:10.1088/1755-1315/1561/1/012018.

[20] A.A. Farzin, Y.-M. Pun, P. Braun, A. Lesage-Landry, Y. Diouane, and I. Shames. Min-Max Optimisation for Nonconvex-Nonconcave Functions Using a Random Zeroth-Order Extragradient Algorithm. *Transactions on Machine Learning Research*, 2025. URL: https://openreview.net/forum?id=1bxY1uAXyr.

[21] S. Goldberg and Z. Liu. The diffusion of networking technologies. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1577–1594, New Orleans, Louisiana, 2013. Society for Industrial and Applied Mathematics. URL: https://dl.acm.org/doi/10.5555/2627817.2627930.

[22] S. Gratton, C.W. Royer, L.N. Vicente, and Z. Zhang. Direct search based on probabilistic feasible descent for bound and linearly constrained problems. *Computational Optimization and Applications*, 72(3):525–559, 2019. doi:10.1007/s10589-019-00062-4.

[23] E. Hallé-Hannan, C. Audet, Y. Diouane, S. Le Digabel, and P. Saves. A distance for mixed-variable and hierarchical domains with meta variables. *Neurocomputing*, 653:131208, 2025. doi:10.1016/j.neucom.2025.131208.

[24] W. Hare, G. Jarry-Bolduc, S. Kerleau, and C.W. Royer. Using orthogonally structured positive bases for constructing positive k-spanning sets with cosine measure guarantees. *Linear Algebra and its Applications*, 680:183–207, 2024. doi:10.1016/j.laa.2023.10.006.

[25] W. Hare and S. Sun. On the computation of the cosine measure in high dimensions. Technical Report 2506.19723, arXiv, 2025. doi:10.48550/arXiv.2506.19723.

[26] C. Huang and Z. Zhang. Non-convergence Analysis of Probabilistic Direct Search. Technical Report 2606.01320, arXiv, 2026. doi:10.48550/arXiv.2606.01320.

[27] Z. Huang, J. Ryga, A. Ogilvy, A. Ardestani-Jaafari, A. Jirasek, and W. Hare. A multiobjective multifidelity optimization approach for the design of scanners for use in optical computed tomography imaging. *Engineering Optimization*, pages 1–20, 2026. doi:10.1080/0305215X.2026.2671223.

[28] S. Kerleau. *Graphical and geometrical perspectives on positive spanning sets with applications to derivative-free optimization*. PhD thesis, Université Paris sciences et lettres, 2025. doi:10.70675/c4c6caf2z4727z4fdaz98ddze6a134622849.

[29] T.G. Kolda, R.M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45(3):385–482, 2003. doi:10.1137/S003614450242889.

[30] J. Larson, M. Menickelly, and E. Toler. Birkhoff interpolation models for optimization with some available derivatives. Technical Report 2605.28503, arXiv, 2026. doi:10.48550/arXiv.2605.28503.

[31] R.M. Lewis, A. Shepherd, and V. Torczon. Implementing Generating Set Search Methods for Linearly Constrained Minimization. *SIAM Journal on Scientific Computing*, 29(6):2507–2530, 2007. doi:10.1137/050635432.

[32] F. Li, I. Kocar, and A. Lesage-Landry. Practical Solutions to Volt-Var Optimization Under Uncertainty via Blackbox Optimization. *IEEE Transactions on Power Delivery*, 41(3):1180–1191, 2026. doi:10.1109/TPWRD.2026.3658012.

[33] V. Mai, Q.H. Pham, A. Favrel, J.-P. Gauthier, and M. Gagnon. Active learning-based optimization of hydroelectric turbine startup to minimize fatigue damage. *Renewable Energy*, 256:124088, 2026. doi:10.1016/j.renene.2025.124088.

[34] F. Marini, M. Porcelli, and E. Riccietti. A multilevel stochastic regularized first-order method with application to finite sum minimization. Technical Report 2412.11630, arXiv, 2026. To appear in *Mathematics of Computation*. doi:10.48550/arXiv.2412.11630.

[35] D. Oh and N. Sahinidis. A large-scale benchmarking of deterministic and stochastic derivative-free optimization algorithms, 2026. doi:10.2172/3031240.

[36] A. Palmieri, F. Rinaldi, and S. Shashaani. Extrapolation-based Direct Search for Nonsmooth Stochastic Zeroth-Order Optimization. Technical Report 35069, Optimization Online, 2026. URL: https://optimization-online.org/?p=35069.

[37] V. Picheny. A Stepwise uncertainty reduction approach to constrained global optimization. In S. Kaski and J. Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 787–795, Reykjavik, Iceland, 2014. URL: https://proceedings.mlr.press/v33/picheny14.

[38] M. Porcelli and Ph.L. Toint. BFO, A Trainable Derivative-free Brute Force Optimizer for Nonlinear Bound-constrained Optimization and Equilibrium Computations with Continuous and Discrete Variables. *ACM Transactions on Mathematical Software*, 44(1):6:1–6:25, 2017. doi:10.1145/3085592.

[39] L. Roberts and C.W. Royer. Polling Set Construction and Worst-Case Complexity for Direct Search under Polyhedral Convex Constraints. Technical Report 2605.27814, arXiv, 2026. doi:10.48550/arXiv.2605.27814.

[40] M. Scheffer, J. Bascompte, W.A. Brock, V. Brovkin, S.R. Carpenter, V. Dakos, H. Held, E.H. van Nes, M. Rietkerk, and G. Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009. doi:10.1038/nature08227.

[41] K. Scheinberg and Z. Xiong. Function-free Optimization via Comparison Oracles. Technical Report 2604.26867, arXiv, 2026. doi:10.48550/arXiv.2604.26867.

[42] E.J. Silva and A.L. Custódio. An inexact restoration direct multisearch filter approach to multiobjective constrained derivative-free optimization. *Optimization Methods and Software*, 40(2):406–432, 2025. doi:10.1080/10556788.2024.2412646.

[43] C. Soldati, S. Le Digabel, and A. Lesage-Landry. Blackbox Optimization for Loss Minimization in Power Distribution Networks using Feeder Reconfiguration. Technical report, CIGRE, Palais des Congrès de Montréal, Canada, 2025. URL: https://www.e-cigre.org/publications/detail/symp-mon-2025-symposium-montreal.html.