

RESEARCH CENTRE

Inria Saclay Centre

2024

ACTIVITY REPORT

Project-Team

RANDOPT

Randomized Optimization

IN COLLABORATION WITH: Centre de Mathématiques Appliquées (CMAP)

DOMAIN

**Applied Mathematics, Computation and
Simulation**

THEME

**Optimization, machine learning and
statistical methods**

Inria

Contents

Project-Team RANDOPT	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	2
2.1 Scientific Context	2
3 Research program	5
3.1 Developing Novel Theoretical Frameworks for Analyzing and Designing Adaptive Stochastic Algorithms	5
3.2 Algorithmic developments	7
3.2.1 Constrained optimization	7
3.2.2 Large-scale Optimization	7
3.2.3 Multiobjective Optimization	8
3.2.4 Expensive Optimization	8
3.3 Setting novel standards in scientific experimentation and benchmarking	8
4 Application domains	9
5 Social and environmental responsibility	10
5.1 Footprint of research activities	10
5.2 Impact of research results	10
6 Highlights of the year	10
6.1 Awards	10
7 New software, platforms, open data	10
7.1 Code Development	10
7.1.1 The <code>pycma</code> package	10
7.1.2 The <code>COCO</code> package	11
7.2 New software	11
7.2.1 <code>COCO</code>	11
7.2.2 <code>CMA-ES</code>	11
7.2.3 <code>COMO-CMA-ES</code>	12
7.2.4 <code>MOarchiving</code>	12
8 New results	13
8.1 Linear convergence proof of <code>CMA-ES</code> and beyond	13
8.2 Mixed-integer optimization	13
8.3 Multi-objective optimization	14
8.4 Benchmarking: methodology and the Comparing Continuous Optimizers Platform (<code>COCO</code>)	14
9 Bilateral contracts and grants with industry	15
9.1 Bilateral contracts with industry	15
10 Partnerships and cooperations	15
10.1 European initiatives	15
10.1.1 Other european programs/initiatives	15
10.2 National initiatives	16
11 Dissemination	16
11.1 Promoting scientific activities	16
11.1.1 Scientific events: organisation	16
11.1.2 Scientific events: selection	16
11.1.3 Journal	17
11.1.4 Invited talks	17

11.1.5 Tutorials	17
11.1.6 Research administration	17
11.2 Teaching - Supervision - Juries	18
11.2.1 Teaching	18
11.2.2 Juries	18
11.2.3 Supervision	18
11.2.4 Internal or external Inria responsibilities	18
11.3 Popularization	19
11.3.1 Participation in Live events	19
12 Scientific production	19
12.1 Major publications	19
12.2 Publications of the year	20
12.3 Cited publications	21

Project-Team RANDOPT

Creation of the Project-Team: 2019 January 01

Keywords

Computer sciences and digital sciences

A6.2.1. – Numerical analysis of PDE and ODE

A6.2.2. – Numerical probability

A6.2.6. – Optimization

A8.2. – Optimization

A8.9. – Performance evaluation

Other research topics and application domains

B4.3. – Renewable energy production

B5.2. – Design and manufacturing

1 Team members, visitors, external collaborators

Research Scientists

- Anne Auger [Team leader, INRIA, Senior Researcher]
- Dimo Brockhoff [INRIA, Researcher]
- Nikolaus Hansen [INRIA, Senior Researcher]

PhD Students

- Mohamed Gharafi [INRIA]
- Oskar Lucien Girardin [INRIA, from May 2024]
- Armand Gissler [INRIA, from Oct 2024]
- Armand Gissler [Ecole Polytechnique, until Sep 2024]
- Tristan Marty [THALES]

Interns and Apprentices

- Nikita Fediashin [INRIA, Intern, from May 2024 until Sep 2024]
- Oskar Lucien Girardin [INRIA, Intern, until Feb 2024]
- Hanwen Xu [INRIA, Intern, from Sep 2024]

Administrative Assistant

- Amandine Sainsard [INRIA, from Mar 2024]

2 Overall objectives

2.1 Scientific Context

Critical problems of the 21st century like the search for highly energy efficient or even carbon-neutral, and cost-efficient systems, or the design of new molecules against extensively drug-resistant bacteria crucially rely on the resolution of challenging numerical optimization problems. Such problems typically depend on noisy experimental data or involve complex numerical simulations where *derivatives are not useful or not available* and the function is considered as a *black-box*.

Many of those optimization problems are in essence *multiobjective*—one needs to optimize simultaneously several conflicting objectives like minimizing the cost of an energy network and maximizing its reliability—and most of the *challenging* black-box problems are *non-convex* and *non-smooth* and they combine difficulties related to ill-conditioning, non-separability, and ruggedness (a term that characterizes functions that can be non-smooth but also noisy or multi-modal). Additionally, the objective function can be expensive to evaluate, that is, one function evaluation can take several minutes to hours (it can involve for instance a CFD simulation).

In this context, the use of randomness combined with proper adaptive mechanisms that notably satisfy certain invariance properties (affine invariance, invariance to monotonic transformations) has proven to be one key component for the design of robust global numerical optimization algorithms [47, 33].

The field of adaptive stochastic optimization algorithms has witnessed some important progress over the past 15 years. On the one hand, subdomains like medium-scale unconstrained optimization may be considered as “solved” (particularly, the CMA-ES algorithm, an instance of *Evolution Strategy* (ES)

algorithms, stands out as state-of-the-art method) and considerably better standards have been established in the way benchmarking and experimentation are performed. On the other hand, multiobjective population-based stochastic algorithms became the method of choice to address multiobjective problems when a set of some best possible compromises is sought after. In all cases, the resulting algorithms have been naturally transferred to industry (the CMA-ES algorithm is now regularly used in companies such as Bosch, Total, ALSTOM, . . .) or to other academic domains where difficult problems need to be solved such as physics, biology [52], geoscience [40], or robotics [43]).

ES algorithms also attracted quite some attention in Machine Learning with the OpenAI article *Evolution Strategies as a Scalable Alternative to Reinforcement Learning*. It is shown that the training time for difficult reinforcement learning benchmarks could be reduced from 1 day (with standard RL approaches) to 1 hour using ES [49].¹ Already ten years ago, another impressive application of CMA-ES, how “Computer Sim Teaches Itself To Walk Upright” (published at the conference SIGGRAPH Asia 2013) was presented in the [press in the UK](#).

Several of these important advances around adaptive stochastic optimization algorithms rely to a great extent on works initiated or achieved by the founding members of RandOpt, particularly related to the CMA-ES algorithm and to the Comparing Continuous Optimizer (COCO) benchmarking platform.

Yet, the field of adaptive stochastic algorithms for black-box optimization is relatively young compared to the “classical optimization” field that includes convex and gradient-based optimization. For instance, the state-of-the-art algorithms for unconstrained gradient based optimization like quasi-Newton methods (e.g. the BFGS method) date from the 1970s [31] while the stochastic derivative-free counterpart, CMA-ES dates from the early 2000s [32]. Consequently, in some subdomains with *important practical demands*, not even the most fundamental and basic questions are answered:

- This is the case of *constrained* optimization where one needs to find a solution $x^* \in \mathbb{R}^n$ minimizing a numerical function $\min_{x \in \mathbb{R}^n} f(x)$ while respecting a number of constraints m typically formulated as $g_i(x^*) \leq 0$ for $i = 1, \dots, m$. Only somewhat recently, the fundamental requirement of linear convergence², as in the unconstrained case, has been clearly stated [23].
- In multiobjective optimization, most of the research so far has been focusing on *how to select candidate solutions from one iteration to the next one*. The difficult question of how to *generate* effectively new solutions is not yet answered in a proper way and we know today that simply applying operators from single-objective optimization may not be effective with the current best selection strategies. As a comparison, in the single-objective case, the question of selection of candidate solutions was already solved in the 1980s and 15 more years were needed to solve the trickier question of an effective adaptive strategy to generate new solutions.
- With the current demand to solve larger and larger optimization problems (e.g. in the domain of deep learning), optimization algorithms that scale linearly (in terms of internal complexity, memory and number of function evaluations to reach an ϵ -ball around the optimum) with the problem dimension are nowadays in increasing demand. Not long ago, first proposals of how to reduce the quadratic scaling of CMA-ES have been made without a clear view of what can be achieved in the best case *in practice*. These later variants apply to optimization problems with thousands of variables. The question of designing randomized algorithms capable to handle problems with one or two orders of magnitude more variables effectively and efficiently is still largely open.
- For expensive optimization, standard methods are so called Bayesian optimization (BO) algorithms often based on Gaussian processes. Commonly used examples of BO algorithms are EGO [37],

¹The key behind such an improvement is the parallelization of the algorithm (on thousands of CPUs) that is done in such a way that the communication between the different workers is reduced to only exchanging a vector of permutation of small length (typically less than 100) containing the ranking of candidate solutions on the function to be optimized. In contrast, parallelization of backpropagation requires to exchange the gradient vector of the size of the problem (of the order of 10^6). This reduced communication time is a decisive factor for the impressive speedup.

²In optimization, linear convergence for an algorithm whose estimate of the optimum x^* of f at iteration t is denoted x_t , refers to a convergence where after a certain time (usually once the initialization is forgotten) the following typically holds: $\|x_{t+1} - x^*\| \leq c \|x_t - x^*\|$ where $c < 1$. This type of convergence is also called geometric. In the case of stochastic algorithms, there exist different definitions of linear convergence (depending on whether we consider the expectation of the sequence or we want a statement that holds with high probability) not strictly equivalent but that always translate the idea that the distance to the optimum at iteration $t + 1$ is a fraction of the distance to the optimum at iteration t .

SMAC [35], Spearmint [50], or TPE [26] which are implemented in different libraries. Yet, our experience with a popular method like EGO is that many important aspects to come up with a good implementation rely on insider knowledge and are not standard across implementations. Two EGO implementations can differ for example in how they perform the initial design, which bandwidth for the Gaussian kernel is used, or which strategy is taken to optimize the expected improvement.

Additionally, the **development of stochastic adaptive methods for black-box optimization has been mainly driven by heuristics and practice**—rather than a general theoretical framework—validated by intensive computational simulations. Undoubtedly, **this has been an asset as the scope of possibilities for design was not restricted by mathematical frameworks** for proving convergence. In effect, powerful stochastic adaptive algorithms for **unconstrained optimization** like the CMA-ES algorithm emerged from this approach. At the same time, naturally, **theory strongly lags behind practice**. For instance, the striking performances of CMA-ES empirically observed contrast with how little is theoretically proven on the method. This situation is clearly not satisfactory. On the one hand, theory generally lifts performance assessment from an empirical level to a conceptual one, rendering results independent from the problem instances where they have been obtained. On the other hand, theory typically provides insights that change perspectives on some algorithm components. Also theoretical guarantees generally increase the trust in the reliability of a method and facilitate the task to make it accepted by wider communities.

Finally, as discussed above, the development of novel black-box algorithms strongly relies on scientific experimentation, and it is quite difficult to conduct proper and meaningful experimental analysis. This is well known for more than two decades now and summarized in this quote from Johnson in 1996

“the field of experimental analysis is fraught with pitfalls. In many ways, the implementation of an algorithm is the easy part. The hard part is successfully using that implementation to produce meaningful and valuable (and publishable!) research results.” [36]

Since then, quite some progress has been made to set better standards in conducting scientific experiments and benchmarking. Yet, some domains still suffer from poor benchmarking standards and from the generic problem of the lack of reproducibility of results. For instance, in multiobjective optimization, it is (still) not rare to see comparisons between algorithms made by solely visually inspecting Pareto fronts after a fixed budget. In Bayesian optimization, good performance seems often to be due to insider knowledge not always well described in papers.

In the context of black-box numerical optimization previously described, the scientific positioning of the RandOpt ream is at the intersection between theory, algorithm design, and applications. Our vision is that the field of stochastic black-box optimization should reach the same level of maturity than gradient-based convex mathematical optimization. This entails major algorithmic developments for constrained, multiobjective and large-scale black-box optimization and major theoretical developments for analyzing current methods including the state-of-the-art CMA-ES.

The specificity in black-box optimization is that methods are intended to solve problems characterized by “*non-properties*”—*non-linear, non-convex, non-smooth, non-Lipschitz*. This contrasts with gradient-based optimization and poses on the one hand some challenges when developing theoretical frameworks but also makes it compulsory to complement theory with empirical investigations.

On the practical side, our ultimate goal is to provide software that is suitable for researchers and industry that need to solve practical optimization problems. We see theory also as a means for this end (rather than only an end in itself) and we also firmly believe that parameter tuning is part of the algorithm designer’s task.

This shapes, on the one hand, four main scientific objectives for our team:

1. **develop novel theoretical frameworks** for guiding (a) the design of novel black-box methods and (b) their analysis, allowing to
2. provide **proofs of key features** of stochastic adaptive algorithms including the state-of-the-art method CMA-ES: linear convergence and learning of second order information.

3. develop **stochastic numerical black-box algorithms** following a **principled design** in domains with a strong practical need for much better methods namely **constrained, multiobjective, large-scale and expensive optimization**. Implement the methods such that they are easy to use. And finally, to
4. **set new standards in scientific experimentation, performance assessment and benchmarking** both for optimization on continuous or combinatorial search spaces. This should allow in particular to advance the state of **reproducibility of results of scientific papers** in optimization.

On the other hand, the above motivates our objectives with respect to dissemination and transfer:

1. develop software packages that people can directly use to solve their problems. This means having carefully thought out interfaces, generically applicable setting of parameters and termination conditions, proper treatment of numerical errors, catching properly various exceptions, etc.;
2. have direct collaborations with industrials;
3. publish our results both in applied mathematics and computer science bridging the gap between very often disjoint communities.

3 Research program

The lines of research we intend to pursue is organized along four axes namely developing novel theoretical framework, developing novel algorithms, setting novel standards in scientific experimentation and benchmarking and applications.

3.1 Developing Novel Theoretical Frameworks for Analyzing and Designing Adaptive Stochastic Algorithms

Stochastic black-box algorithms typically optimize **non-convex, non-smooth functions**. This is possible because the algorithms rely on weak mathematical properties of the underlying functions: the algorithms do not use the derivatives—hence the function does not need to be differentiable—and, additionally, often do not use the exact function value but instead how the objective function ranks candidate solutions (such methods are sometimes called function-value-free). (To illustrate a comparison-based update, consider an algorithm that samples λ (with λ an even integer) candidate solutions from a multivariate normal distribution. Let x_1, \dots, x_λ in \mathbb{R}^n denote those λ candidate solutions at a given iteration. The solutions are evaluated on the function f to be minimized and ranked from the best to the worse:

$$f(x_{1:\lambda}) \leq \dots \leq f(x_{\lambda:\lambda}) .$$

In the previous equation $i:\lambda$ denotes the index of the sampled solution associated to the i -th best solution. The new mean of the Gaussian vector from which new solutions will be sampled at the next iteration can be updated as

$$m \leftarrow \frac{4}{\lambda} \sum_{i=1}^{\lambda/4} x_{i:\lambda} .$$

The previous update moves the mean towards the $\lambda/2$ best solutions. Yet the update is only based on the ranking of the candidate solutions such that the update is the same if f is optimized or $g \circ f$ where $g: \text{Im}(f) \rightarrow \mathbb{R}$ is strictly increasing. Consequently, such algorithms are invariant with respect to strictly increasing transformations of the objective function. This entails that they are robust and their performances generalize well.)

Additionally, adaptive stochastic optimization algorithms typically have a **complex state space** which encodes the parameters of a probability distribution (e.g. mean and covariance matrix of a Gaussian vector) and other state vectors. This state-space is a **manifold**. While the algorithms are Markov chains, the complexity of the state-space makes that **standard Markov chain theory tools do not directly apply**. The same holds with tools stemming from stochastic approximation theory or Ordinary Differential Equation (ODE) theory where it is usually assumed that the underlying ODE (obtained by proper averaging

and limit for learning rate to zero) has its critical points inside the search space. In contrast, in the cases we are interested in, the **critical points of the ODEs are at the boundary of the domain**.

Last, since we aim at developing theory that on the one hand allows to analyze the main properties of state-of-the-art methods and on the other hand is useful for algorithm design, we need to be careful not to use simplifications that would allow a proof to be done but would not capture the important properties of the algorithms. With that respect one tricky point is to develop **theory that accounts for invariance properties**.

To face those specific challenges, we need to develop novel theoretical frameworks exploiting invariance properties and accounting for peculiar state-spaces. Those frameworks should allow researchers to analyze one of the core properties of adaptive stochastic methods, namely **linear convergence** on the widest possible class of functions.

We are planning to approach the question of linear convergence from three different complementary angles, using three different frameworks:

- the Markov chain framework where the convergence derives from the analysis of the stability of a normalized Markov chain existing on scaling-invariant functions for translation and scale-invariant algorithms [25]. This framework allows for a fine analysis where the exact convergence rate can be given as an implicit function of the invariant measure of the normalized Markov chain. Yet it requires the objective function to be scaling-invariant. The stability analysis can be particularly tricky as the Markov chain that needs to be studied writes as $\Phi_{t+1} = F(\Phi_t, W_{t+1})$ where $\{W_t : t > 0\}$ are independent identically distributed and F is typically discontinuous because the algorithms studied are comparison-based. This implies that practical tools for analyzing a standard property like irreducibility, that rely on investigating the stability of underlying deterministic control models [44, Chapter 7], cannot be used. Additionally, the construction of a drift to prove ergodicity is particularly delicate when the state space includes a (normalized) covariance matrix as it is the case for analyzing the CMA-ES algorithm.
- The stochastic approximation or ODE framework. Those are standard techniques to prove the convergence of stochastic algorithms when an algorithm can be expressed as a stochastic approximation of the solution of a mean field ODE [28, 27, 41]. What is specific and induces difficulties for the algorithms we aim at analyzing is the **non-standard state-space** since the ODE variables correspond to the state-variables of the algorithm (e.g. $\mathbb{R}^n \times \mathbb{R}_{>0}$ for step-size adaptive algorithms, $\mathbb{R}^n \times \mathbb{R}_{>0} \times S_{++}^n$ where S_{++}^n denotes the set of positive definite matrices if a covariance matrix is additionally adapted). Consequently, the ODE can have many critical points at the boundary of its definition domain (e.g. all points corresponding to $\sigma_t = 0$ are critical points of the ODE) which is not typical. Also we aim at proving **linear convergence**, for that it is crucial that the learning rate does not decrease to zero which is non-standard in ODE method.
- The direct framework where we construct a global Lyapunov function for the original algorithm from which we deduce bounds on the hitting time to reach an ϵ -ball of the optimum. For this framework as for the ODE framework, we expect that the class of functions where we can prove linear convergence are composite of $g \circ f$ where f is differentiable and $g : \text{Im}(f) \rightarrow \mathbb{R}$ is strictly increasing and that we can show convergence to a local minimum.

We expect those frameworks to be complementary in the sense that the assumptions required are different. Typically, the ODE framework should allow for proofs under the assumptions that learning rates are small enough while it is not needed for the Markov chain framework. Hence this latter framework captures better the real dynamics of the algorithm, yet under the assumption of scaling-invariance of the objective functions. Also, we expect some overlap in terms of function classes that can be studied by the different frameworks (typically convex-quadratic functions should be encompassed in the three frameworks). By studying the different frameworks in parallel, we expect to gain synergies and possibly understand what is the most promising approach for solving the holy grail question of the linear convergence of CMA-ES. We foresee for instance that similar approaches like the use of Foster-Lyapunov drift conditions are needed in all the frameworks and that intuition can be gained on how to establish the conditions from one framework to another one.

3.2 Algorithmic developments

We are planning on developing algorithms in the subdomains with strong practical demand for better methods of constrained, multiobjective, large-scale and expensive optimization.

Many of the algorithm developments, we propose, rely on the CMA-ES method. While this seems to restrict our possibilities, we want to emphasize that CMA-ES became a *family of methods* over the years that nowadays include various techniques and developments from the literature to handle non-standard optimization problems (noisy, large-scale, ...). The core idea of all CMA-ES variants—namely the mechanism to adapt a Gaussian distribution—has furthermore been shown to derive naturally from first principles with only minimal assumptions in the context of derivative-free black-box stochastic optimization [47, 33]. This is a strong justification for relying on the CMA-ES premises while new developments naturally include new techniques typically borrowed from other fields. While CMA-ES is now a full family of methods, for visibility reasons, we continue to refer often to “the CMA-ES algorithm”.

3.2.1 Constrained optimization

Many (real-world) optimization problems have constraints related to technical feasibility, cost, etc. Constraints are classically handled in the black-box setting either via rejection of solutions violating the constraints—which can be quite costly and even lead to quasi-infinite loops—or by penalization with respect to the distance to the feasible domain (if this information can be extracted) or with respect to the constraint function value [29]. However, the penalization coefficient is a sensitive parameter that needs to be adapted in order to achieve a robust and general method [30]. Yet, **the question of how to handle properly constraints is largely unsolved**. Previous constraints handling for CMA-ES were ad-hoc techniques driven by many heuristics [30]. Also, only somewhat recently it was pointed out that **linear convergence properties should be preserved** when addressing constraint problems [23].

Promising approaches though, rely on using augmented Lagrangians [23, 24]. The augmented Lagrangian, here, is the objective function optimized by the algorithm. Yet, it depends on coefficients that are adapted online. The adaptation of those coefficients is the difficult part: the algorithm should be stable and the adaptation efficient. We believe that the theoretical frameworks developed (particularly the Markov chain framework) will be useful to understand how to design the adaptation mechanisms. Additionally, the question of invariance will also be at the core of the design of the methods: augmented Lagrangian approaches break the invariance to monotonic transformation of the objective functions, yet understanding the maximal invariance that can be achieved seems to be an important step towards understanding what adaptation rules should satisfy.

3.2.2 Large-scale Optimization

In the large-scale setting, we are interested to optimize problems with the order of 10^3 to 10^4 variables. For one to two orders of magnitude more variables, we will talk about a “very large-scale” setting.

In this context, algorithms with a quadratic scaling (internal and in terms of number of function evaluations needed to optimize the problem) cannot be afforded. In CMA-ES-type algorithms, we typically need to restrict the model of the covariance matrix to have only a linear number of parameters to learn such that the algorithms scale linearly in terms of internal complexity, memory and number of function evaluations to solve the problem. The main challenge is thus to have rich enough models for which we can efficiently design proper adaptation mechanisms. Some first large-scale variants of CMA-ES have been derived. They include the online adaptation of the complexity of the model [22, 21]. Yet, the type of Hessian matrices they can learn is restricted and not fully satisfactory. Different restricted families of distributions are conceivable and it is an open question which can be effectively learned and which are the most promising in practice.

Another direction, we want to pursue, is exploring the use of large-scale variants of CMA-ES to solve reinforcement learning problems [49].

Last, we are interested to investigate the very-large-scale setting. One approach consists in doing optimization in subspaces. This entails the efficient identification of relevant spaces and the restriction of the optimization to those subspaces.

3.2.3 Multiobjective Optimization

Multiobjective optimization, i.e., the simultaneous optimization of multiple objective functions, differs from single-objective optimization in particular in its optimization goal. Instead of aiming at converging to the solution with the best possible function value, in multiobjective optimization, a set of solutions³ is sought. This set, called Pareto-set, contains all trade-off solutions in the sense of Pareto-optimality—no solution exists that is better in *all* objectives than a Pareto-optimal one. Because converging towards a set differs from converging to a single solution, it is no surprise that we might lose many good convergence properties if we directly apply search operators from single-objective methods. However, this is what has typically been done so far in the literature. Indeed, most of the research in stochastic algorithms for multiobjective optimization focused instead on the so called selection part, that decides which solutions should be kept during the optimization—a question that can be considered as solved for many years in the case of single-objective stochastic adaptive methods.

We therefore aim at rethinking search operators and adaptive mechanisms to improve existing methods. We expect that we can obtain orders of magnitude better convergence rates for certain problem types if we choose the right search operators. We typically see two angles of attack: On the one hand, we will study methods based on scalarizing functions that transform the multiobjective problem into a set of single-objective problems. Those single-objective problems can then be solved with state-of-the-art single-objective algorithms. Classical methods for multiobjective optimization fall into this category, but they all solve multiple single-objective problems subsequently (from scratch) instead of dynamically changing the scalarizing function during the search. On the other hand, we will improve on currently available population-based methods such as the first multiobjective versions of the CMA-ES. Here, research is needed on an even more fundamental level such as trying to understand success probabilities observed during an optimization run or how we can introduce non-elitist selection (the state of the art in single-objective stochastic adaptive algorithms) to increase robustness regarding noisy evaluations or multi-modality. The challenge here, compared to single-objective algorithms, is that the quality of a solution is not anymore independent from other sampled solutions, but can potentially depend on all known solutions (in the case of three or more objective functions), resulting in a more noisy evaluation as the relatively simple function-value-based ranking within single-objective optimizers.

3.2.4 Expensive Optimization

In the so-called expensive optimization scenario, a single function evaluation might take several minutes or even hours in a practical setting. Hence, the available budget in terms of number of function evaluation calls to find a solution is very limited in practice. To tackle such expensive optimization problems, it is needed to exploit the first few function evaluations in the best way. To this end, typical methods couple the learning of a surrogate (or meta-model) of the expensive objective function with traditional optimization algorithms.

In the context of expensive optimization and CMA-ES, which usually shows its full potential when the number n of variables is not too small (say larger than 3) and if the number of available function evaluations is about $100n$ or larger, several research directions emerge. The two main possibilities to integrate meta-models into the search with CMA-ES type algorithms are (i) the successive injection of the minimum of a learned meta-model at each time step into the learning of CMA-ES's covariance matrix and (ii) the use of a meta-model to predict the internal ranking of solutions. While for the latter, first results exist, the former idea is entirely unexplored for now. In both cases, a fundamental question is which type of meta-model (linear, quadratic, Gaussian Process, ...) is the best choice for a given number of function evaluations (as low as one or two function evaluations) and at which time the type of the meta-model shall be switched.

3.3 Setting novel standards in scientific experimentation and benchmarking

Numerical experimentation is needed as a complement to theory to test novel ideas, hypotheses, the stability of an algorithm, and/or to obtain quantitative estimates. Optimally, theory and experimentation go hand in hand, jointly guiding the understanding of the mechanisms underlying optimization algorithms.

³Often, this set forms a manifold of dimension one smaller than the number of objectives.

Though performing numerical experimentation on optimization algorithms is crucial and a common task, it is non-trivial and easy to fall in (common) pitfalls as stated by J. N. Hooker in his seminal paper [34].

In the RandOpt team we aim at raising the standards for both scientific experimentation and benchmarking.

On the experimentation aspect, we are convinced that there is common ground over how scientific experimentation should be done across many (sub-)domains of optimization, in particular with respect to the visualization of results, testing extreme scenarios (parameter settings, initial conditions, etc.), how to conduct understandable and small experiments, how to account for invariance properties, performing scaling up experiments and so forth. We therefore want to formalize and generalize these ideas in order to make them known to the entire optimization community with the final aim that they become standards for experimental research.

Extensive numerical benchmarking, on the other hand, is a compulsory task for evaluating and comparing the performance of algorithms. It puts algorithms to a standardized test and allows to make recommendations which algorithms should be used preferably in practice. To ease this part of optimization research, we have been developing the Comparing Continuous Optimizers platform (COCO) since 2007 which allows to automatize the tedious task of benchmarking. It is a game changer in the sense that the freed time can now be spent on the scientific part of algorithm design (instead of implementing the experiments, visualization, etc.) and it opened novel perspectives in algorithm testing. COCO implements a thorough, well-documented methodology that is based on the above mentioned general principles for scientific experimentation.

Also due to the freely available data from 350+ algorithms benchmarked with the platform, COCO became a quasi-standard for single-objective, noiseless optimization benchmarking. It is therefore natural to extend the reach of COCO towards other subdomains (particularly constrained optimization, many-objective optimization) which can benefit greatly from an automated benchmarking methodology and standardized tests without (much) effort. This entails particularly the design of novel test suites and rethinking the methodology for measuring performance and more generally evaluating the algorithms. Particularly challenging is the design of scalable non-trivial testbeds for constrained optimization where one can still control where the solutions lies. Other optimization problem types, we are targeting are expensive problems (and the Bayesian optimization community in particular), optimization problems in machine learning (for example parameter tuning in reinforcement learning), and the collection of real-world problems from industry.

Another aspect of our future research on benchmarking is to investigate the large amounts of benchmarking data, we collected with COCO during the years. Extracting information about the influence of algorithms on the best performing portfolio, clustering algorithms of similar performance, or the automated detection of anomalies in terms of good/bad behavior of algorithms on a subset of the functions or dimensions are some of the ideas here.

Last, we want to expand the focus of COCO from automatized (large) benchmarking experiments towards everyday experimentation, for example by allowing the user to visually investigate algorithm internals on the fly or by simplifying the set up of algorithm parameter influence studies.

4 Application domains

Applications of black-box algorithms occur in various domains. Industry but also researchers in other academic domains have a great need to apply black-box algorithms on a daily basis. Generally, we do not target a specific application domain and are interested in black-box applications stemming from various origins. This is to us intrinsic to the nature of the methods we develop that are general purpose algorithms. Hence our strategy with respect to applications can be considered as opportunistic and our main selection criteria when approached by colleagues who want to develop a collaboration around an application is whether we find the application interesting and valuable: that means the application brings new challenges and/or gives us the opportunity to work on topics we already intended to work on, and it brings, in our judgement, an advancement to society in the application domain.

The concrete applications related to industrial collaborations we are currently dealing with are:

- With Thales for the theses of Konstantinos Varelas, Paul Dufossé, and Tristan Marty (DGA-CIFRE

theses). They investigate more specifically the development of large-scale variants of CMA-ES, constrained-handling for CMA-ES, and the handling of discrete variables within CMA-ES respectively.

- With Storengy, a subsidiary of the ENGIE group, specialized in gas storage for the theses of Cheikh Touré and Mohamed Gharafi. Different multiobjective applications are considered in this context but the primary motivation of Storengy is to get at their disposal a better multiobjective variant of CMA-ES which is the main objective of the developments within the theses.

5 Social and environmental responsibility

5.1 Footprint of research activities

We are concerned about CO₂ footprint and discourage oversea conferences when far away. Since the situation with respect to Covid went back to normal with respect to travelling, we have been dedicated to travel less than in the past and attend some conferences online.

5.2 Impact of research results

We develop general purpose optimization methods that apply in difficult optimization contexts where little is required on the function to be optimized. Application domains include optimization and design of renewable systems and climate change.

Our main method CMA-ES is transferred and widely used. The code stemming from the team is frequently downloaded (see Section 7). Among the usage of our method and our code, we find naturally problems in the domain of energy to capture carbon dioxide [45, 42, 46], solar energy [38, 39], or wind-thermal power systems [48].

Those publications witness the impact of our research results with respect to research questions and engineering design related to climate change and renewable energy.

6 Highlights of the year

One of the most significant achievements of the year for our research team has been the theoretical proof of the linear convergence of CMA-ES, alongside its ability to learn second-order information through the covariance matrix. This breakthrough—which we have been pursuing since more than a decade—directly addresses a key challenge we set out to tackle within the RANDOPT team. Given the widespread impact of CMA-ES, with over 70 million downloads of its two main source codes, this result marks a major milestone in the field of optimization and beyond (see 8.1)

6.1 Awards

Anne Auger received the [SIGEVO outstanding contribution award 2024](#).

7 New software, platforms, open data

7.1 Code Development

7.1.1 The `pycma` package

The `pycma` package had two major releases during 2024. A refactoring of the options and parameters code was undertaken, the interfaces for noise handling and coordinate scaling were revised and a new integer handling was released.

7.1.2 The COCO package

The COCO project was entirely restructured and separated into two major subprojects, the C and Python package `cocoex` and the Python package `cocopp`. The noisy bbob testbed implementation was released in the `cocoex` package. The navigation functionality in the displayed results of `cocopp` were significantly improved, see also Section 8.4.

7.2 New software

7.2.1 COCO

Name: COmparing Continuous Optimizers

Keywords: Benchmarking, Numerical optimization, Black-box optimization, Stochastic optimization

Scientific Description: COmparing Continuous Optimisers (COCO) is a tool for benchmarking algorithms for black-box optimisation. COCO facilitates systematic experimentation in the field of continuous optimization. COCO provides: (1) an experimental framework for testing the algorithms, (2) post-processing facilities for generating publication quality figures and tables, including the easy integration of data from benchmarking experiments of 350+ algorithm variants, (3) LaTeX templates for scientific articles and HTML overview pages which present the figures and tables.

The COCO software is composed of two parts: (i) an interface available in different programming languages (C/C++, Java, Matlab/Octave, Python, external support for R) which allows to run and log experiments on several function test suites (unbounded noisy and noiseless single-objective functions, unbounded noiseless multiobjective problems, mixed-integer problems, constrained problems) and (ii) a Python tool for generating figures and tables that can be looked at in every web browser and that can be used in the provided LaTeX templates to write scientific papers.

Functional Description: The COCO platform aims at supporting the numerical benchmarking of black-box optimization algorithms in continuous domains. Benchmarking is a vital part of algorithm engineering and a necessary path to recommend algorithms for practical applications. The COCO platform releases algorithm developers and practitioners alike from (re-)writing test functions, logging, and plotting facilities by providing an easy-to-handle interface in several programming languages. The COCO platform has been developed since 2007 and has been used extensively within the “Blackbox Optimization Benchmarking (BBOB)” workshop series since 2009. Overall, 350+ algorithms and algorithm variants by contributors from all over the world have been benchmarked on the platform’s supported test suites so far. The most recent extensions has been towards constrained problems.

URL: <https://github.com/numbbo/coco>

Contact: Dimo Brockhoff

Participants: Anne Auger, Asma Atamna, Dejan Tusar, Dimo Brockhoff, Marc Schoenauer, Nikolaus Hansen, Ouassim Ait Elhara, Raymond Ros, Tea Tusar, Thanh-Do Tran, Umut Batu, Konstantinos Varelas

Partners: Charles University Prague, Jozef Stefan Institute (JSI), Cologne University of Applied Sciences

7.2.2 CMA-ES

Name: Covariance Matrix Adaptation Evolution Strategy

Keywords: Numerical optimization, Black-box optimization, Stochastic optimization

Scientific Description: The CMA-ES is considered as state-of-the-art in evolutionary computation and has been adopted as one of the standard tools for continuous optimisation in many (probably hundreds of) research labs and industrial environments around the world. The CMA-ES is typically

applied to unconstrained or bound-constrained optimization problems and search space dimension between three and a few hundred. Recent versions can also handle nonlinear constraints. The method should be applied, if derivative based methods, e.g. quasi-Newton BFGS or conjugate gradient, (supposedly) fail due to a rugged search landscape, e.g. discontinuities, sharp bends or ridges, noise, local optima, outliers. If second order derivative based methods are successful, they are usually much faster than the CMA-ES: on purely convex-quadratic functions, BFGS (Matlabs function `fminunc`) is typically faster by a factor of about ten (in number of objective function evaluations assuming that gradients are not available) and on the most simple quadratic functions by a factor of about 30.

Functional Description: The CMA-ES is an evolutionary algorithm for difficult non-linear non-convex black-box optimisation problems in continuous domain.

URL: <https://cma-es.github.io/>

Contact: Nikolaus Hansen

Participant: Nikolaus Hansen

7.2.3 COMO-CMA-ES

Name: Comma Multi-Objective Covariance Matrix Adaptation Evolution Strategy

Keywords: Black-box optimization, Global optimization, Multi-objective optimisation

Scientific Description: The CMA-ES is considered as state-of-the-art in evolutionary computation and has been adopted as one of the standard tools for continuous optimisation in many (probably hundreds of) research labs and industrial environments around the world. The CMA-ES is typically applied to unconstrained or bounded constraint optimization problems, and search space dimensions between three and a hundred. COMO-CMA-ES is a multi-objective optimization algorithm based on the standard CMA-ES using the Uncrowded Hypervolume Improvement within the so-called Sofomore framework.

Functional Description: The COMO-CMA-ES is an evolutionary algorithm for difficult non-linear non-convex black-box optimisation problems with several (two) objectives in continuous domain.

URL: <https://github.com/CMA-ES/pycomocma>

Contact: Nikolaus Hansen

7.2.4 MOarchiving

Name: Multiobjective Optimization Archiving Module

Keywords: Mathematical Optimization, Multi-objective optimisation

Scientific Description: Multi-objective optimization relies on the maintenance of a set of non-dominated (and hence incomparable) solutions. Performance indicator computations and in particular the computation of the hypervolume indicator is based on this solution set. The hypervolume computation and the update of the set of non-dominated solutions are generally time critical operations. The module computes the bi-objective hypervolume in linear time and updates the non-dominated solution set in logarithmic time.

Functional Description: The module implements a bi-objective non-dominated archive using a Python list as parent class. The main functionality is heavily based on the `bisect` module. The class provides easy and fast access to the overall hypervolume, the contributing hypervolume of each element, and to the uncrowded hypervolume improvement of any given point in objective space.

URL: <https://github.com/CMA-ES/moarchiving>

Contact: Nikolaus Hansen

8 New results

8.1 Linear convergence proof of CMA-ES and beyond

Participants: Anne Auger, Nikolaus Hansen, Armand Gissler.

One of the key theoretical objectives pursued by the RANDOPT team is to establish a rigorous proof of the linear convergence of CMA-ES while demonstrating that it learns second-order information. This fundamental question has remained open for over 20 years and is often regarded as the "holy grail" by researchers in the field. Armand Gissler, who defended his PhD in December 2024, focused on this topic [16]. The main difficulties for achieving the proof could be solved in the context of his PhD.

The proof builds on a methodology we previously developed for analyzing step-size adaptive evolution strategies. We had to resolve key theoretical challenges for achieving the convergence of CMA-ES. The complete proof involves multiple milestones, with results presented across several publications (two of which are still being finalized):

- Asymptotic estimations of the eigenvalue and eigenvectors of a symmetric matrix perturbed by the sum of rank one matrices were obtained [13]. The estimates obtained are crucial for proving the stability of underlying Markov chains of CMA-ES when the covariance matrix has a large condition number.
- The extension of tools to prove the irreducibility of Markov chains living on a manifold via analyzing the stability of a deterministic control model were needed to be able to apply them to CMA-ES. The ensuing work is presented in [18].
- Those tools were then applied to prove the irreducibility of the different normalized Markov chains associated to the different variants of CMA-ES [19].
- Two other publications are in preparation. The first establishes drift conditions for the geometric ergodicity of underlying Markov chains. The second uses stability and invariance to complete the proof using the different milestones.

In another paper, currently under review for the ACM-GECCO conference, we introduce a rank-based surrogate-assisted variant of CMA-ES. Unlike previous methods that employ rank information as constraints to train an SVM classifier, our approach employs a linear-quadratic regression on the ranks. While this approach is invariant under any monotonous transformation of objective function values and also outperforms the original CMA-ES in numerical experiments with a few exceptions, it falls short to entirely meet the performance levels of the surrogate-based lq-CMA-ES (which is not invariant under monotonous transformations of the objective function). To address this, we propose a variant that handles together two alternative surrogates, one based on the ranks and one based on the original function values. Although this variant sacrifices strict invariance, it gains in robustness and achieves performance comparable to, or even exceeding, lq-CMA-ES on transformed problems. This work is achieved in the context of the Action Exploratoire LearnToOpt.

8.2 Mixed-integer optimization

We developed LB+IC-CMA-ES, a novel variant of CMA-ES designed to handle mixed-integer optimization problems. The algorithm incorporates two key mechanisms for managing integer variables: (i) a lower bound (LB) on the variance of integer components to maintain exploration and (ii) integer centering (IC), which adjusts integer variables toward the center of their domain based on their values. We conducted a thorough evaluation of the different algorithmic variants resulting from these modifications using the BBOB mixed-integer benchmark suite. The performance of LB+IC-CMA-ES was compared against the recently introduced CMA-ES with margin, providing insights into its effectiveness in mixed-integer settings [14]. The algorithm was applied to optical problem in the context of the PhD thesis of Tristan Marty [15].

8.3 Multi-objective optimization

Participants: Anne Auger, Dimo Brockhoff, Mohamed Gharafi, Nikolaus Hansen, **external collaborators:** Rodolphe Le Riche (CNRS LIMOS), Tea Tušar, Jordan Cork (Jozef Stefan Institute).

A central theme for the team is the design, analysis, and benchmarking of multiobjective optimization algorithms. In 2024, we have progressed on the following aspects.

In the context of our collaboration with the company Storengy for the PhD thesis of Mohamed Gharafi, we realized that our algorithm COMO-CMA-ES, a multiobjective variant of CMA-ES [51], shows slower and slower convergence speed in numerical experiments when the number of solutions in the sought Pareto-set approximation increases. Additional experiments showed that this is inherent to all algorithms, based on the optimization of the hypervolume indicator or its variants (like for COMO-CMA-ES).

Together with our Bachelor student, Nguyen Vu, we proposed a new so-called continuation method for bi-objective problems, which aims at following the Pareto set (in case it is continuous) from one objective function's optimum to the optimum of the second objective function (and back). The algorithm is based on a two-step, single-objective formulation of the original bi-objective problem, trying to minimize first the distance to the Pareto set (via minimization of the multiobjective gradient) and secondly minimize alternatingly one of the original objective functions. Using the default CMA-ES as single-objective optimizer, we could achieve a working algorithm for the case of continuous Pareto sets. Numerical experiments on the COCO platform shows that the relatively simple algorithm is competitive with other, more complicated algorithms like the unbounded-population size version of MO-CMA-ES for problems with continuous Pareto sets when the budget is relatively small (smaller than 1000 times dimension). The corresponding paper is still under review for the ACM-GECCO conference.

With the goal to extend our COCO platform (see also the next section) towards constrained multiobjective problems, we investigated theoretically the Pareto sets of m -objective problems with spherical objective functions and convex feasible sets. Our main result shows that the Pareto set of the constrained problem can be constructed by projecting the Pareto set of the unconstrained problem onto the convex feasible set. This paper, as well, is still under review for the ACM-GECCO conference.

8.4 Benchmarking: methodology and the Comparing Continuous Optimizers Platform (COCO)

Participants: Anne Auger, Dimo Brockhoff, Nikolaus Hansen, **external collaborators:** Tobias Glasmachers (Ruhr-Universität Bochum), Jakub Kudela (Brno University of Technology), Olaf Mersmann (TH Köln), Tea Tušar (Jozef Stefan Institute), Vanessa Volz (Centrum Wiskunde & Informatica).

Benchmarking is an important task in optimization in order to assess and compare the performance of algorithms as well as to motivate the design of better solvers. We are leading the benchmarking of derivative free solvers in the context of difficult problems: we have been developing methodologies and testbeds as well as assembled this into a platform automatizing the benchmarking process. This is a continuing effort that we are pursuing in the team.

The COCO platform, developed at Inria since 2007, aims at automatizing numerical benchmarking experiments and the visual presentation of their results. The platform consists of an experimental part to generate benchmarking data (in various programming languages) and a postprocessing module (in Python), see Figure 1. At the interface between the two, we provide data sets from numerical experiments of 300+ algorithms and algorithm variants from various fields (quasi-Newton, derivative-free optimization, evolutionary computing, Bayesian optimization) and for various problem characteristics (noiseless/noisy optimization, single-/multi-objective optimization, continuous/mixed-integer, ...).

We have been using the platform in the past to initiate workshop papers during the ACM-GECCO conference as well as to collect algorithm data sets from the entire optimization community (350+ so far

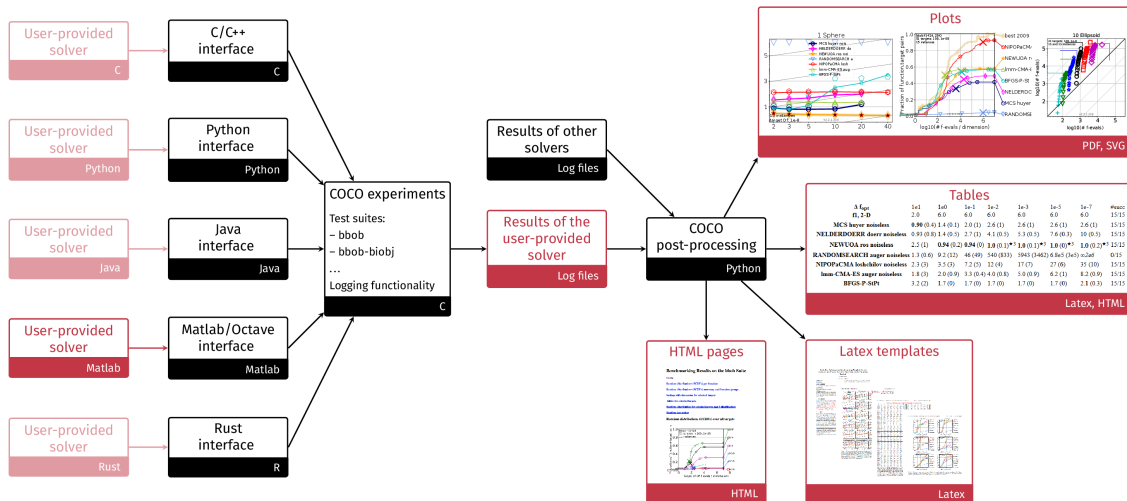


Figure 1: Structural overview of the COCO platform. COCO provides all black parts while users only have to connect their solver to the COCO interface in the language of interest, here for instance Matlab, and to decide on the test suite the solver should run on. The other red components show the output of the experiments (number of function evaluations to reach certain target precisions) and their post-processing and are automatically generated.

over the different test suites). The next of those workshops will be held at the ACM-GECCO conference in Malaga, Spain in summer 2025.

In November 2024, we held our **second COCO Documentation and Coding Sprint** at Schloss Dagstuhl, Germany with a group of seven developers and focused on a complete re-design of the **COCO webpage**, new functionalities in the postprocessing module, a new scenario of outliers (or frozen noise) with varying noise level, and fixing of other important bugs. We also made a large step forward to further automatize our continuous integration tests. Over the four-and-a-half days of the COCO sprint, we can report 250+ commits with more than 25,000 lines of code touched.

9 Bilateral contracts and grants with industry

Participants: Anne Auger, Dimo Brockhoff, Nikolaus Hansen, Tristan Marty, Mohamed Garafi.

9.1 Bilateral contracts with industry

- Contract with the company Storengy funding the PhD thesis of Mohamed Gharafi in the context of the CIROQUO project (2021–2024), for the latest report, see [17]
- Contract with Thales for the CIFRE PhD thesis of Tristan Marty (2023–2026)

10 Partnerships and cooperations

10.1 European initiatives

10.1.1 Other european programs/initiatives

COST Action ROAR-NET

Participants: Anne Auger, Dimo Brockhoff, Nikolaus Hansen.

Title: COST Action ROAR-NET (“Randomised Optimisation Algorithms Research Network”, CA22137)

Partner Institution(s): 35 partner countries with 381 researchers in total

Date/Duration: 2023–2027

Additional info/keywords: The permanent Randopt members are involved in the two working groups “Mixed Continuous and Discrete Optimisation” (WG2) and “Single- and Multiobjective Optimisation” (WG3).

10.2 National initiatives

CIROQUO

Participants: Dimo Brockhoff, Mohamed Gharafi, Nikolaus Hansen.

Title: CIROQUO (“Consortium Industriel de Recherche en Optimisation et QUantification d’incertitudes pour les données Onéreuses”)

Partner Institution(s): six other academic and five industrial partners, including Storengy

Date/Duration: 2021–2024

Additional info/keywords: Randopt is involved in the context of the PhD thesis of Mohamed Gharafi which is financed by the company Storengy.

11 Dissemination

Participants: Anne Auger, Dimo Brockhoff, Nikolaus Hansen.

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

- A. Auger, D. Brockhoff, and N. Hansen: organization of the Blackbox Optimization Benchmarking Workshop during the ACM-GECCO-2025 conference, together with Tobias Glasmachers (Bochum, Germany), Olaf Mersmann (Brühl, Germany), and Tea Tušar (Ljubljana, Slovenia)
- [2nd COCO Code and Documentation Sprint](#), Schloss Dagstuhl, Germany, November 24–29, 2024

Member of the organizing committees

- A. Auger is co-organizer of the 2024 Dagstuhl seminar on [Theory of Randomized Search Heuristics](#).

11.1.2 Scientific events: selection

Chair of conference program committees

- D. Brockhoff: co-track chair of the Evolutionary Multiobjective Optimization track (EMO), ACM-GECCO-2024 conference, together with T. Ray
- A. Auger, scientific committee of the FGS (French-German-Spanish) Conference on Optimization 2024

Member of the conference program committees

- RandOpt permanent members are frequent reviewer and in the program committee of the ACM GECCO and PPSN conference.

Reviewer

- D. Brockhoff: PPSN 2024, EMO 2025, ACM-GECCO 2025, ACM-FOGA 2025
- A. Auger: PPSN 2024, ACM-GECCO 2024

11.1.3 Journal

Member of the editorial boards

- D. Brockhoff: associate editor for ACM Transactions on Evolutionary Learning and Optimization (TELO), since 2019
- N. Hansen: editorial board member of *Evolution Computation Journal*, MIT (2023 Impact Factor: 4.6).
- A. Auger: editorial board member of *Evolution Computation Journal*, MIT (2023 Impact Factor: 4.6).

Reviewer - reviewing activities

- A. Auger: reviewer for ECJ, IEEE Trans on ECJ
- D. Brockhoff: reviewer for ACM Transactions on Evolutionary Learning and Optimization

11.1.4 Invited talks

- A. Auger, invited speaker at the Journées SMAI MODE, March 2024
- A. Auger, invited speaker at the 2nd Derivative-free Optimization Symposium (DFOS), Padova, June 2024
- N. Hansen, invited speaker at the 2nd Derivative-free Optimization Symposium (DFOS), Padova, June 2024
- D. Brockhoff, invited speaker, Dagstuhl seminar on "Theory of Randomized search heuristics", July 2024

11.1.5 Tutorials

- N. Hansen: Tutorial *CMA-ES* at the [PPSN Parallel Problem Solving from Nature conference](#), Hagenberg, Austria.

11.1.6 Research administration

- A. Auger: member of the **management committee** of the European COST action *Randomised Optimisation Algorithms Research Network*, **head** of the working group on mixed-integer optimization till October 2024.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Master: A. Auger, “Derivative-free Optimization”, 22.5h ETD, niveau M2 (Optimization Master of Paris-Saclay)
- Master: D. Brockhoff, “Algorithms and Complexity”, 36h ETD, niveau M1 (joint MSc with ESSEC “Data Sciences & Business Analytics”), CentraleSupélec, France
- Master: D. Brockhoff, “Advanced Optimization”, 36h ETD, niveau M2 (joint MSc with ESSEC “Data Sciences & Business Analytics”), CentraleSupélec, France
- Master: D. Brockhoff, “Introduction to Optimization”, 36h ETD, niveau M1 (joint MSc with ESSEC “Data Sciences & Business Analytics”),
- Bachelor: A. Auger, “Convex Optimization and Optimal Control”, Bachelor of Ecole Polytechnique, 3rd year.

11.2.2 Juries

- A. Auger, reviewer and PhD jury of **Achille Jacquemond**, Ecole Centrale Lyon (defended in October 2024)
- A. Auger, head of the PhD jury of **Thomas Guilmeau**, January 2025.
- N. Hansen, reviewer and PhD jury member of **Paul Templier**, ISAE-SUPAERO Université de Toulouse, April 2024.

11.2.3 Supervision

PhD students

- Mohamed Gharafi, supervised by D. Brockhoff, N. Hansen, and R. Le Riche (CNRS, Ecole des Mines Saint-Etienne)
- Oskar Lucien Girardin, supervised by D. Brockhoff and N. Hansen, since May 2024
- Armand Gissler, supervised by A. Auger and N. Hansen
- Tristan Marty, supervised by A. Auger and N. Hansen

Master's and Bachelor's students

- Nikita Fediashin, supervised by A. Auger, from May 2024 until Sep 2024
- Oskar Lucien Girardin, M2 internship, CentraleSupélec/ESSEC, supervised by D. Brockhoff and N. Hansen, until February 2024
- Hanwen Xu, M2 internship, CentraleSupélec/ESSEC, supervised by D. Brockhoff, from Sep 2024
- Nguyen Vu, Bachelor project, Ecole Polytechnique, supervised by D. Brockhoff

11.2.4 Internal or external Inria responsibilities

- D. Brockhoff: member of the CUMI at Inria Saclay, since October 2023
- D. Brockhoff: correspondent AMIES for the CMAP, since January 2023
- D. Brockhoff: member of the CDT at Inria Saclay, since February 2019

11.3 Popularization

- Welcomed Mathias Onimus (High school student, 1st year) for an internship of 2 weeks.

11.3.1 Participation in Live events

- A. Auger: participation in the Science Day at Ecole Polytechnique (Inria booth).
- A. Auger, scientific presentation for high school students in the context of the two weeks internship that 1st year high school students have to do (at Inria Saclay).

12 Scientific production

12.1 Major publications

- [1] Y. Akimoto, A. Auger, T. Glasmachers and D. Morinaga. ‘Global Linear Convergence of Evolution Strategies on More Than Smooth Strongly Convex Functions’. In: *SIAM Journal on Optimization* (June 2022). URL: <https://hal.inria.fr/hal-02941429>.
- [2] Y. Akimoto, A. Auger and N. Hansen. ‘An ODE Method to Prove the Geometric Convergence of Adaptive Stochastic Algorithms’. In: *Stochastic Processes and their Applications* 145 (2022), pp. 269–307. DOI: [10.1016/j.spa.2021.12.005](https://doi.org/10.1016/j.spa.2021.12.005). URL: <https://hal.inria.fr/hal-01926472>.
- [3] Y. Akimoto and N. Hansen. ‘Diagonal Acceleration for Covariance Matrix Adaptation Evolution Strategies’. In: *Evolutionary Computation* 28.3 (2020), pp. 405–435. DOI: [10.1162/evco_a_00260](https://doi.org/10.1162/evco_a_00260). URL: <https://hal.inria.fr/hal-01995373>.
- [4] A. Auger and N. Hansen. ‘A SIGEVO impact award for a paper arising from the COCO platform’. In: *ACM SIGEVOlution* 13.4 (Jan. 2021), pp. 1–11. DOI: [10.1145/3447929.3447930](https://doi.org/10.1145/3447929.3447930). URL: <https://hal.inria.fr/hal-03132960>.
- [5] D. Brockhoff, A. Auger, N. Hansen and T. Tušar. ‘Using Well-Understood Single-Objective Functions in Multiobjective Black-Box Optimization Test Suites’. In: *Evolutionary Computation* (2022). DOI: [10.1162/evco_a_00298](https://doi.org/10.1162/evco_a_00298). URL: <https://hal.inria.fr/hal-01296987>.
- [6] A. Chotard and A. Auger. ‘Verifiable Conditions for the Irreducibility and Aperiodicity of Markov Chains by Analyzing Underlying Deterministic Models’. In: *Bernoulli* 25.1 (Dec. 2018), pp. 112–147. DOI: [10.3150/17-BEJ970](https://doi.org/10.3150/17-BEJ970). URL: <https://hal.inria.fr/hal-01222222>.
- [7] N. Hansen. ‘A Global Surrogate Assisted CMA-ES’. In: *GECCO 2019 - The Genetic and Evolutionary Computation Conference*. ACM, Prague, Czech Republic, July 2019, pp. 664–672. DOI: [10.1145/3321707.3321842](https://doi.org/10.1145/3321707.3321842). URL: <https://hal.inria.fr/hal-02143961>.
- [8] N. Hansen, A. Auger, R. Ros, O. Mersmann, T. Tušar and D. Brockhoff. ‘COCO: A Platform for Comparing Continuous Optimizers in a Black-Box Setting’. In: *Optimization Methods and Software* 36.1 (2020). ArXiv e-prints, arXiv:1603.08785, pp. 114–144. DOI: [10.1080/10556788.2020.1808977](https://doi.org/10.1080/10556788.2020.1808977). URL: <https://hal.inria.fr/hal-01294124>.
- [9] Y. Ollivier, L. Arnold, A. Auger and N. Hansen. ‘Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles’. In: *Journal of Machine Learning Research* 18.18 (2017), pp. 1–65. URL: <https://hal.inria.fr/hal-01515898>.
- [10] C. Touré, A. Auger and N. Hansen. ‘Global linear convergence of Evolution Strategies with recombination on scaling-invariant functions’. In: *Journal of Global Optimization* (2022). DOI: [10.1007/s10898-022-01249-6](https://doi.org/10.1007/s10898-022-01249-6). URL: <https://hal.science/hal-03286037>.
- [11] C. Touré, A. Gissler, A. Auger and N. Hansen. ‘Scaling-invariant functions versus positively homogeneous functions’. In: *Journal of Optimization Theory and Applications* (Sept. 2021). URL: <https://hal.inria.fr/hal-03104436>.

- [12] C. Touré, N. Hansen, A. Auger and D. Brockhoff. ‘Uncrowded Hypervolume Improvement: COMO-CMA-ES and the Sofomore framework’. In: *GECCO 2019 - The Genetic and Evolutionary Computation Conference*. Part of this research has been conducted in the context of a research collaboration between Storengy and Inria. Prague, Czech Republic, July 2019. DOI: [10.1145/3321707.3321852](https://doi.org/10.1145/3321707.3321852). URL: <https://hal.inria.fr/hal-02103694>.

12.2 Publications of the year

International journals

- [13] A. Gissler, A. Auger and N. Hansen. ‘Asymptotic estimations of a perturbed symmetric eigenproblem’. In: *Applied Mathematics Letters* 150 (2024), p. 108951. DOI: [10.1016/j.aml.2023.108951](https://doi.org/10.1016/j.aml.2023.108951). URL: <https://inria.hal.science/hal-04386103> (cit. on p. 13).

International peer-reviewed conferences

- [14] T. Marty, N. Hansen, A. Auger, Y. Semet and S. Héron. ‘LB+IC-CMA-ES: Two Simple Modifications of CMA-ES to Handle Mixed-Integer Problems’. In: *Parallel Problem Solving from Nature — PPSN XVIII, 18th International Conference*. Vol. 15149. Lecture Notes in Computer Science. Hagenberg, Austria: Springer Nature Switzerland, 7th Sept. 2024, pp. 284–299. DOI: [10.1007/978-3-031-70068-2_18](https://doi.org/10.1007/978-3-031-70068-2_18). URL: <https://inria.hal.science/hal-04740412> (cit. on p. 13).
- [15] T. Marty, S. Héron, Y. Semet, A. Auger and N. Hansen. ‘Glass optimization for optical design: CMA-ES optimizer with integer handling’. In: *Proceedings Volume PC13023, Computational Optics 2024*. SPIE Optical Systems Design. Strasbourg, France: SPIE, 7th Apr. 2024, p. 7. DOI: [10.1117/12.3016937](https://doi.org/10.1117/12.3016937). URL: <https://inria.hal.science/hal-04937349> (cit. on p. 13).

Doctoral dissertations and habilitation theses

- [16] A. Gissler. ‘Linear convergence of evolution strategies with covariance matrix adaptation’. École polytechnique, 16th Dec. 2024. URL: <https://hal.science/tel-04935192> (cit. on p. 13).

Reports & preprints

- [17] C. Blanchet-Scalliet, C. Helbert, D. Sinoquet, M. M. Munoz Zuniga, R. Le Riche, D. Rullière, C. Prieur, O. Zahm, J. Garnier, D. Brockhoff, O. Roustant, F. Bachoc, L. Pronzato, Y. Richet, J. Rohmer, F. Huguet, D. Gaudrie, A. Marrel, B. Kerleguer, T. Perrillat-Bottonet and C. Durantin. *Activity report ciroquo research & industry consortium*. Ecole Centrale de Lyon; Mines Saint-Etienne; Université Toulouse 3 (Paul Sabatier); Stellantis France; BRGM (Bureau de recherches géologiques et minières); CEA; IFP Energies Nouvelles; Institut de Radioprotection et de Sûreté Nucléaire; Storengy; INRIA; CNRS, 2024, pp. 1–11. URL: <https://hal.science/hal-04661116> (cit. on p. 15).
- [18] A. Gissler, A. Durmus and A. Auger. *On the irreducibility and convergence of a class of nonsmooth nonlinear state-space models on manifolds*. 6th Feb. 2024. URL: <https://inria.hal.science/hal-04713677> (cit. on p. 13).
- [19] A. Gissler, S.-C. Wolf, A. Auger and N. Hansen. *Irreducibility of nonsmooth state-space models with an application to CMA-ES*. 29th Sept. 2024. URL: <https://inria.hal.science/hal-04713675> (cit. on p. 13).

Other scientific publications

- [20] N. Hansen. *CMA-ES: Covariance Matrix Adaptation Evolution Strategy - Tutorial*. Hagenberg, Austria, 14th Sept. 2024. URL: <https://inria.hal.science/hal-04709819>.

12.3 Cited publications

- [21] Y. Akimoto and N. Hansen. ‘Online model selection for restricted covariance matrix adaptation’. In: *International Conference on Parallel Problem Solving from Nature*. Springer, 2016, pp. 3–13 (cit. on p. 7).
- [22] Y. Akimoto and N. Hansen. ‘Projection-based restricted covariance matrix adaptation for high dimension’. In: *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference*. ACM, 2016, pp. 197–204 (cit. on p. 7).
- [23] D. V. Arnold and J. Porter. ‘Towards an Augmented Lagrangian Constraint Handling Approach for the (1 + 1)-ES’. In: *Genetic and Evolutionary Computation Conference*. ACM Press, 2015, pp. 249–256 (cit. on pp. 3, 7).
- [24] A. Atamna, A. Auger and N. Hansen. ‘Linearly Convergent Evolution Strategies via Augmented Lagrangian Constraint Handling’. In: *Foundation of Genetic Algorithms (FOGA)*. 2017 (cit. on p. 7).
- [25] A. Auger and N. Hansen. ‘Linear Convergence of Comparison-based Step-size Adaptive Randomized Search via Stability of Markov Chains’. In: *SIAM Journal on Optimization* 26.3 (2016), pp. 1589–1624 (cit. on p. 6).
- [26] J. Bergstra, R. Bardenet, Y. Bengio and B. Kégl. ‘Algorithms for Hyper-Parameter Optimization’. In: *Neural Information Processing Systems (NIPS 2011)*. 2011. URL: <https://hal.inria.fr/hal-00642998/file/draft1.pdf> (cit. on p. 4).
- [27] V. Borkar and S. Meyn. ‘The O.D.E. Method for Convergence of Stochastic Approximation and Reinforcement Learning’. In: *SIAM Journal on Control and Optimization* 38.2 (Jan. 2000) (cit. on p. 6).
- [28] V. S. Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Cambridge University Press, 2008 (cit. on p. 6).
- [29] C. A. Coello Coello. ‘Constraint-handling techniques used with evolutionary algorithms’. In: *Proceedings of the 2008 Genetic and Evolutionary Computation Conference*. ACM, 2008, pp. 2445–2466 (cit. on p. 7).
- [30] G. Collange, S. Reynaud and N. Hansen. ‘Covariance Matrix Adaptation Evolution Strategy for Multidisciplinary Optimization of Expendable Launcher Families’. In: *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Proceedings*. 2010 (cit. on p. 7).
- [31] J. E. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Englewood Cliffs, NJ: Prentice-Hall, 1983 (cit. on p. 3).
- [32] N. Hansen and A. Ostermeier. ‘Completely Derandomized Self-Adaptation in Evolution Strategies’. In: *Evolutionary Computation* 9.2 (2001), pp. 159–195 (cit. on p. 3).
- [33] N. Hansen and A. Auger. ‘Principled design of continuous stochastic search: From theory to practice’. In: *Theory and principled methods for the design of metaheuristics*. Springer, 2014, pp. 145–180 (cit. on pp. 2, 7).
- [34] J. N. Hooker. ‘Testing heuristics: We have it all wrong’. In: *Journal of heuristics* 1.1 (1995), pp. 33–42 (cit. on p. 9).
- [35] F. Hutter, H. Hoos and K. Leyton-Brown. ‘An Evaluation of Sequential Model-based Optimization for Expensive Blackbox Functions’. In: *GECCO (Companion) 2013*. Amsterdam, The Netherlands: ACM, 2013, pp. 1209–1216 (cit. on p. 4).
- [36] D. S. Johnson. ‘A theoretician’s guide to the experimental analysis of algorithms’. In: *Data structures, near neighbor searches, and methodology: fifth and sixth DIMACS implementation challenges* 59 (2002), pp. 215–250 (cit. on p. 4).
- [37] D. R. Jones, M. Schonlau and W. J. Welch. ‘Efficient global optimization of expensive black-box functions’. In: *Journal of Global optimization* 13.4 (1998), pp. 455–492 (cit. on p. 3).
- [38] J. H. Kämpf and D. Robinson. ‘A hybrid CMA-ES and HDE optimisation algorithm with application to solar energy potential’. In: *Applied Soft Computing* 9.2 (2009), pp. 738–745 (cit. on p. 10).

- [39] J. H. Kämpf and D. Robinson. ‘Optimisation of building form for solar energy utilisation using constrained evolutionary algorithms’. In: *Energy and Buildings* 42.6 (2010), pp. 807–814 (cit. on p. 10).
- [40] I. Kriest, V. Sauerland, S. Khatiwala, A. Srivastav and A. Oschlies. ‘Calibrating a global three-dimensional biogeochemical ocean model (MOPS-1.0)’. In: *Geoscientific Model Development* 10.1 (2017), p. 127 (cit. on p. 3).
- [41] H. J. Kushner and G. Yin. *Stochastic approximation and recursive algorithms and applications*. Applications of mathematics. New York: Springer, 2003 (cit. on p. 6).
- [42] X. Lu, B. Ganis and M. F. Wheeler. ‘Optimal Design of CO₂ Sequestration with Three-Way Coupling of Flow-Geomechanics Simulations and Evolution Strategy’. In: *SPE Reservoir Simulation Conference*. OnePetro. 2019 (cit. on p. 10).
- [43] P. MacAlpine, S. Barrett, D. Urieli, V. Vu and P. Stone. ‘Design and Optimization of an Omnidirectional Humanoid Walk: A Winning Approach at the RoboCup 2011 3D Simulation Competition’. In: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*. Toronto, Ontario, Canada, July 2012 (cit. on p. 3).
- [44] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. New York: Springer-Verlag, 1993 (cit. on p. 6).
- [45] A. Miyagi, Y. Akimoto and H. Yamamoto. ‘Well placement optimization for carbon dioxide capture and storage via CMA-ES with mixed integer support’. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2018, pp. 1696–1703 (cit. on p. 10).
- [46] A. Miyagi, H. Yamamoto, Y. Akimoto and Z. Xue. ‘Development of a high speed optimization tool for well placement in Geological Carbon dioxide Sequestration’. In: *5th ISRM Young Scholars’ Symposium on Rock Mechanics and International Symposium on Rock Engineering for Innovative Future*. OnePetro. 2019 (cit. on p. 10).
- [47] Y. Ollivier, L. Arnold, A. Auger and N. Hansen. In: *The Journal of Machine Learning Research* 18.1 (2017), pp. 564–628 (cit. on pp. 2, 7).
- [48] S. S. Reddy, B. Panigrahi, R. Kundu, R. Mukherjee and S. Debchoudhury. ‘Energy and spinning reserve scheduling for a wind-thermal power system using CMA-ES with mean learning technique’. In: *International Journal of Electrical Power & Energy Systems* 53 (2013), pp. 113–122 (cit. on p. 10).
- [49] T. Salimans, J. Ho, X. Chen and I. Sutskever. ‘Evolution strategies as a scalable alternative to reinforcement learning’. In: *arXiv preprint arXiv:1703.03864* (2017) (cit. on pp. 3, 7).
- [50] J. Snoek, H. Larochelle and R. P. Adams. ‘Practical bayesian optimization of machine learning algorithms’. In: *Neural Information Processing Systems (NIPS 2012)*. 2012, pp. 2951–2959 (cit. on p. 4).
- [51] C. Touré, N. Hansen, A. Auger and D. Brockhoff. ‘Uncrowded Hypervolume Improvement: COMO-CMA-ES and the Sofomore framework’. In: *GECCO 2019 - The Genetic and Evolutionary Computation Conference*. Part of this research has been conducted in the context of a research collaboration between Storengy and Inria. Prague, Czech Republic, July 2019. DOI: [10.1145/3321707.3321852](https://doi.org/10.1145/3321707.3321852). URL: <https://hal.inria.fr/hal-02103694> (cit. on p. 14).
- [52] J. Uhlenendorf, A. Miermont, T. Delaveau, G. Charvin, F. Fages, S. Bottani, G. Batt and P. Hersen. ‘Long-term model predictive control of gene expression at the population and single-cell levels’. In: *Proceedings of the National Academy of Sciences* 109.35 (2012), pp. 14271–14276 (cit. on p. 3).