Activity Report 2019

Project-Team RANDOPT

Randomized Optimization

IN COLLABORATION WITH: Centre de Mathématiques Appliquées (CMAP)

Table of contents

Creation of the Team: 2016 December 01, updated into Project-Team: 2019 January 01

Keywords:

Computer Science and Digital Science:

A6.2. - Scientific computing, Numerical Analysis & Optimization
A8.2. - Optimization
A8.9. - Performance evaluation

Other Research Topics and Application Domains:

B4.4. - Energy delivery
B9.5.1. - Computer science
B9.5.2. - Mathematics
B9.8. - Reproducibility

1. Team, Visitors, External Collaborators

Research Scientists

Anne Auger [Team leader, Inria, Researcher, HDR]
Dimo Brockhoff [Inria, Researcher]
Nikolaus Hansen [Inria, Senior Researcher, HDR]

PhD Students

Alann Cheral [École polytechnique, PhD Student, from Oct 2019]
Marie Ange Dahito [PSA, PhD Student, granted by CIFRE]
Paul Dufosse [Thales, PhD Student, granted by CIFRE]
Cheikh Saliou Toure [Inria, PhD Student]
Konstantinos Varelas [Thales, PhD Student, granted by CIFRE]
Eugenie Marescaux [Inria, PhD Student, from Sep 2019]

Interns and Apprentices

Julien Bonneville [Inria, from May 2019 until Aug 2019]
Eugenie Marescaux [Inria, from May 2019 until Aug 2019]

Administrative Assistants

Natalia Alves [Inria, Administrative Assistant, until Jan 2019]
Marie Enee [Inria, Administrative Assistant, from Mar 2019]

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 such that *derivatives are not useful or not available* and the function is seen 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, non-smooth* and 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 particularly satisfy several invariance properties (affine invariance, invariance to monotonic transformations) has proven to be one key component for the design of robust global numerical optimization algorithms [35], [25].

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 thought for. 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 [38], geoscience [31], or robotics [33]).

Very recently, ES algorithms 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 [36]. [1] A few 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 those important advances around adaptive stochastic optimization algorithms are relying 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) 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 [24] while the stochastic derivative-free counterpart, CMA-ES dates from the early 2000s [26]. 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 $i = 1, ..., m$. Only recently, the fundamental requirement of linear convergence [2], as in the unconstrained case, has been clearly stated [16].

---

[1] The 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 an important factor for the important speedup.

[2] In 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.

- 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 needed. Only recently, 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 efficiently problems with one or two orders of magnitude more variables 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 [30], SMAC [28], Spearmint [37], or TPE [19] 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 tested. 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." [29]

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.

2.2. Overall Objectives

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 a *non*-property—*non*-convex, *non*-linear, *non*-smooth. 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.

Our ultimate goal is to provide software that is useful for practitioners. We see that theory is a means for this end (rather than an end in itself) and it is also our firm belief that parameter tuning is part of the designer's task.

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

3.1. Introduction

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

3.2. 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 $x_1, ..., x_\lambda$ 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:

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

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

$$m \longleftarrow \frac{1}{\lambda} \sum_{i=1}^{\lambda/2} x_{i:\lambda} \ .$$

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 $g : \text{Im}(f) \to \mathbb{R}$ 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 [18]. 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 [34], 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 [20], [21], [32]. 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 $g : \mathrm{Im}(f) \to \mathbb{R}$ 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.

3.3. 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 [35], [25]. 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".

3.3.1. 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 [22]. However, the penalization coefficient is a sensitive parameter that needs to be adapted in order to achieve a robust and general method [23]. Yet, **the question of how to handle properly constraints is largely unsolved**. The latest constraints handling for CMA-ES is an ad-hoc technique driven by many heuristics [23]. Also, it is particularly only recently that it was pointed out that **linear convergence properties should be preserved** when addressing constraint problems [16].

Promising approaches though, rely on using augmented Lagrangians [16], [17]. 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.

3.3.2. 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 [15], [14]. Yet so far they fail to optimize functions whose Hessian matrix has some small eigenvalues (say around 10^{-4}) some eigenvalues equal to 1 and some very large eigenvalue (say around 10^4), that is functions whose level sets have short and long axis.

Another direction, we want to pursue, is exploring the use of large-scale variants of CMA-ES to solve reinforcement learning problems [36].

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.

3.3.3. 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 solutions [3] 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

---

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

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.

3.3.4. *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.

3.4. 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. 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 [27].

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, statistical tests, 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 300+ 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, see our AESOP project), 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

4.1. 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. We do not target a specific application domain and are interested in possible black-box applications stemming from various origins. This is for us intrinsic to the nature of the methods we develop that are general purpose algorithms. Hence our strategy with respect to applications can be seen as opportunistic and our main selection criteria when approached by colleagues who want to develop a collaboration around an application is whether we judge the application interesting: that is the application brings new challenges and/or gives us the opportunity to work on topics we already intended to work on.

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

- With Thales for the theses of Konstantinos Varelas and Paul Dufossé (DGA-CIFRE theses) related to the design of radars (shape optimization of the wave form). Those theses investigate the development of large-scale variants of CMA-ES and constrained-handling for CMA-ES.

- With Storengy, a subsidiary of Engie specialized in gas storage for the thesis of Cheikh Touré. 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 thesis.

- With PSA in the context of the OpenLab and the thesis of Marie-Ange Dahito for the design of part of a car body.
- With Onera in the context of the thesis of Alann Cheral related to the optimization of the choice of hyperspectral bandwidth.

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. Awards

- Cheikh Touré received the 2nd prize for the student best paper award at the conference EMO 2019 for the paper "On Bi-Objective convex-quadratic problems" by Cheikh Touré, Anne Auger, Dimo Brockhoff, Nikolaus Hansen
- Nikolaus Hansen received the ENUM best paper award at the ACM-GECCO 2019 conference (see https://gecco-2019.sigevo.org/index.html/Best+Paper+Awards) for the paper "A global Surrogate Assisted CMA-ES"

BEST PAPERS AWARDS:

[6]
C. TOURÉ, A. AUGER, D. BROCKHOFF, N. HANSEN. *On Bi-Objective convex-quadratic problems*, in "10th International Conference on Evolutionary Multi-Criterion Optimization", East Lansing, Michigan, United States, March 2019, https://arxiv.org/abs/1812.00289 , https://hal.inria.fr/hal-01942159

[5]
N. HANSEN. *A Global Surrogate Assisted CMA-ES*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, ACM, 2019, pp. 664-672 [*DOI :* 10.1145/3321707.3321842], https://hal.inria.fr/hal-02143961

6. New Software and Platforms

6.1. COCO

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 300+ 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, constrained problems) are provided (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 blackbox 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, 300+ 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 extension has been towards large-scale problems and was used for the BBOB-2019 workshop at the ACM-GECCO conference.

- 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 and Konstantinos Varelas
- Partners: TU Dortmund University - Charles University Prague - Jozef Stefan Institute (JSI)
- Contact: Dimo Brockhoff
- URL: https://github.com/numbbo/coco

6.2. CMA-ES

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 bounded constraint optimization problems, and search space dimensions between three and a hundred. 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 faster than the CMA-ES: on purely convex-quadratic functions, $f(x)=x^THx$, BFGS (Matlabs function fminunc) is typically faster by a factor of about ten (in terms of number of objective function evaluations needed to reach a target function value, assuming that gradients are not available). On the most simple quadratic function $f(x)=\|x\|^2=x^Tx$ BFGS is faster 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.

- Participant: Nikolaus HANSEN
- Contact: Nikolaus HANSEN
- URL: http://cma.gforge.inria.fr/cmaes_sourcecode_page.html

7. New Results

7.1. A multiobjective algorithm framework and the COMO-CMA-ES algorithm

One classical way to solve multiobjective optimization problems is to transform the multiple objective functions into a single one, also known as scalarization, and to solve multiple versions of such scalarizations to achieve multiple trade-off points. Another approach, coming from the field of evolutionary multiobjective optimization is to formulate a single-objective *set* problem via indicators: the goal here is to find the set of solutions (of a given size) that maximizes a certain quality. The hypervolume indicator has been regularly used to define the quality of the solution set because it has favorable theoretical properties.

The "classical" definition of the hypervolume indicator and how it is used in multiobjective solvers, however, has some disadvantages: (i) the resulting single-objective problem is of high dimension and the gradient of the hypervolume indicator is zero in dominated areas of the search space—not giving a solver enough information about where to search for good solutions. In [7], we discussed and visualized these disadvantages and proposed a new quality criterion which is based on the hypervolume indicator but solves the mentioned disadvantages. The implementation of this idea and its combination with the single-objective solver CMA-ES resulted in the COMO-CMA-ES which shows improved performance over several existing multiobjective solvers on a wide range of convex quadratic functions.

7.2. A mixed-integer benchmark testbed for single and multiobjective black-box optimization

We have introduced two suites of mixed-integer benchmark problems to be used for analyzing and comparing black-box optimization algorithms. They contain problems of diverse difficulties that are scalable in the number of decision variables. The bbob-mixint suite is designed by partially discretizing the established BBOB (Black-Box Optimization Benchmarking) problems. The bi-objective problems from the bbob-biobj-mixint suite are, on the other hand, constructed by using the bbob-mixint functions as their separate objectives. We explain the rationale behind our design decisions and show how to use the suites within the COCO (Comparing Continuous Optimizers) platform. Analyzing two chosen functions in more detail, we also provide some unexpected findings about their properties [8].

7.3. A large-scale optimization testbed for the COCO framework

We have finalized a large scale testbed built to model well-known difficulties in continuous optimization and to test the scaling behavior of algorithms. It contrasts with current test suites used for benchmarking solvers.

The test suite contains 24 single-objective functions in continuous domain and extends the well-known single-objective noiseless bbob test suite. A core contribution is to reduce the computational demand of the orthogonal search space transformations, that appear in the bbob test suite, while retaining some desired properties using previously introduced permuted block diagonal orthogonal matrices.

The paper discusses the implementation details, particularly the normalization and scaling to obtain backwards compatibility with the bbob test suite. Additionally, a guide for using the test suite within the COCO platform, as well as a description of the postprocessed output is presented [12].

7.4. Diagonal Acceleration for Covariance Matrix Adaptation Evolution Strategies

In [1], we have introduced an acceleration for the covariance matrix adaptation evolution strategies (CMA-ES) by means of adaptive diagonal decoding (dd-CMA). This diagonal acceleration endows the default CMA-ES with the advantages of separable CMA-ES without inheriting its drawbacks. Technically, we introduce a diagonal matrix that expresses coordinate-wise variances of the sampling distribution. The diagonal matrix can learn a rescaling of the problem in the coordinates within linear number of function evaluations. Diagonal decoding can also exploit separability of the problem, but, crucially, does not compromise the performance on non-separable problems. The latter is accomplished by modulating the learning rate for the diagonal matrix based on the condition number of the underlying correlation matrix. dd-CMA-ES not only combines the advantages of default and separable CMA-ES, but may achieve overadditive speedup: it improves the performance, and even the scaling, of the better of default and separable CMA-ES on classes of non-separable test functions that reflect, arguably, a landscape feature commonly observed in practice. The paper makes two further secondary contributions: we introduce two different approaches to guarantee positive definiteness of the covariance matrix with active CMA, which is valuable in particular with large population size; we revise the default parameter setting in CMA-ES, proposing accelerated settings in particular for large dimension. All our contributions can be viewed as independent improvements of CMA-ES, yet they are also complementary and

can be seamlessly combined. In numerical experiments with dd-CMA-ES up to dimension 5120, we observe remarkable improvements over the original covariance matrix adaptation on functions with coordinate-wise ill-conditioning. The improvement is observed also for large population sizes up to about dimension squared.

7.5. A global surrogate assisted CMA-ES

In the paper [5], we have explored the arguably simplest way to build an effective surrogate fitness model in continuous search spaces. The model complexity is linear or diagonal-quadratic or full quadratic, depending on the number of available data. The model parameters are computed from the Moore-Penrose pseudoinverse. The model is used as a surrogate fitness for CMA-ES if the rank correlation between true fitness and surrogate value of recently sampled data points is high. Otherwise, further samples from the current population are successively added as data to the model. We empirically compare the IPOP scheme of the new model assisted lq-CMA-ES with a variety of previously proposed methods and with a simple portfolio algorithm using SLSQP and CMA-ES. We conclude that a global quadratic model and a simple portfolio algorithm are viable options to enhance CMA-ES. The model building code is available as part of the pycma Python module on Github and PyPI.

7.6. Benchmarking and Understanding Optimizers

Benchmarking is an important task in optimization in order to understand the working principles behind existing solvers, to find out about weaknesses of them and to finally recommend good ones. The COCO platform, developed in the Randopt team since 2007, aims at automatizing these numerical benchmarking experiments and the visual presentation of their results. We regularly use the platform to initiate workshop papers during the ACM-GECCO conference and also held a workshop this year [4].

In this context, several workshop papers have been published by members of the team and we also proposed some extensions of the platform and updated its documentation.

Two papers addressed single-objective unconstrained problems. One paper investigated the impact of the sample volume of a simple random search on the bbob test suite of COCO [2] and the other paper benchmarked all solvers available in the scipy.optimize module of Python [9], re-discovering SLSQP as a very well-performing solver for small budgets.

Two additional papers addressed multiobjective problems in the context of the bbob-biobj test suite of COCO: "Benchmarking Algorithms from the platypus Framework on the Biobjective bbob-biobj Testbed" [3] compared several baseline algorithms from the literature such as NSGA-II, MOEA/D, SPEA2, and IBEA and "Benchmarking MO-CMA-ES and COMO-CMA-ES on the Bi-objective bbob-biobj Testbed" [4] compared our new COMO-CMA-ES solver with its previous version MO-CMA-ES.

As to extensions of the COCO platform, we released new test suites this year as described earlier. For the large-scale test suite of [12], 11 algorithm variants of the CMA-ES and L-BFGS solvers have been compared in the paper "Benchmarking Large Scale Variants of CMA-ES and L-BFGS-B on the bbob-largescale Testbed" [10].

Overall, we collected 54 new algorithm data sets within the COCO platform in 2019—the highest number in a single year since the release of COCO.

Finally, we updated our documentation on the biobjective test suite, we introduced in 2019. The corresponding journal paper "Using Well-Understood Single-Objective Functions in Multiobjective Black-Box Optimization Test Suites" [11] is now under revision for the Evolutionary Computation journal.

---

[4]See numbbo.github.io/workshops/BBOB-2019/

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

- Contract with the company Storengy partially funding the PhD thesis of Cheikh Touré (2017–2020)
- Contract with Thales in the context of the CIFRE PhD thesis of Konstantinos Varelas (2017–2020)
- Contract with PSA in the context of the CIFRE PhD thesis of Marie-Ange Dahito (2019–2022)
- Pending contract for the thesis of Paul Dufossé with Thales (2020–2022)

9. Partnerships and Cooperations

9.1. Regional Initiatives

- PGMO/FMJH project "AESOP: Algorithms for Expensive Simulation-Based Optimization", 7kEUR, 2017–2019

9.2. National Initiatives

9.2.1. ANR

- ANR project "Big Multiobjective Optimization (BigMO)", Dimo Brockhoff participates in this project through the Inria team BONUS in Lille (2017–2020)

9.3. International Initiatives

9.3.1. Inria International Partners

9.3.1.1. Informal International Partners

- Youhei Akimoto, Tsukuba University, Japan
- Tobias Glasmachers, Ruhr University, Bochum, Germany
- Tea Tušar, Jozef Stefan Institute, Ljubljana, Slovenia

9.4. International Research Visitors

9.4.1. Visits to International Teams

9.4.1.1. Research Stays Abroad

Anne Auger and Dimo Brockhoff visited Tea Tušar (Jozef Stefan Institute, Slovenia) for two weeks in April 2019

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events: Organisation

10.1.1.1. General Chair, Scientific Chair

- Anne Auger: general chair of the ACM-GECCO 2019 conference that took place in July in Prague and welcomed ca. 700 participants.

10.1.1.2. Member of the Organizing Committees

- Dimo Brockhoff: co-organizer of the Lorentz Center Workhop "MACODA: Many Criteria Optimization and Decision Analysis", September 2019, Leiden, The Netherlands, 42 participants
- Anne Auger, Dimo Brockhoff, Nikolaus Hansen, and Konstantinos Varelas: co-organizers of the Black-Box Optimization Benchmarking workshop (BBOB) at the ACM-GECCO 2019 conference.
- Anne Auger: Organizer of an **invited session** on *Recent Advances on Randomized Derivative Free Optimization* at the EURO conference, Dublin.

10.1.2. Scientific Events: Selection

10.1.2.1. Reviewer

- Dimo Brockhoff and Nikolaus Hansen: ACM-GECCO 2019

10.1.3. Dagstuhl seminar invitations

- Anne Auger, Dimo Brockhoff and Nikolaus Hansen invited at the Dagstuhl Seminar 19431 on "Theory of Randomized Optimization Heuristics", October 2019
- Dimo Brockhoff invited at the Dagstuhl Seminar 20031 on "Scalability in Multiobjective Optimization", January 2020

10.1.4. Journal

10.1.4.1. Member of the Editorial Boards

- Anne Auger, Dimo Brockhoff and Nikolaus Hansen: Associate Editor of the ACM Transactions on Evolutionary Learning and Optimization
- Anne Auger and Nikolaus Hansen: Associate Editor of the Evolutionary Computation Journal
- Anne Auger is guest editor of an Algorithmica special issue of papers selected from the ACM-GECCO'2018 theory track
- Anne Auger is guest editor of the IEEE Transactions on Evolutionary Computation special issue on Theoretical Foundations of Evolutionary Computation

10.1.4.2. Reviewer - Reviewing Activities

- Dimo Brockhoff: reviewed papers for IEEE Transactions on Evolutionary Computation (IEEETEC), European Journal of Operational Research (EJOR), Evolutionary Computation, and Journal of Global Optimization

10.1.5. Invited Talks

- Dimo Brockhoff: "Quality Indicator Maximization in Multiobjective Optimization Via Single-Objective Solvers: Unflattened Hypervolume Improvement in the Sofomore Framework", Jozef Stefan Institute, Ljubljana, Slovenia, April 2019
- Nikolaus Hansen: "How to Evolve Gradient Descent into Evolution Strategies and CMA-ES", Journée de Rentrée du CMAP, Ecole polytechnique, Paris, October 2019

- Nikolaus Hansen: "How to Evolve Gradient Descent into Evolution Strategies and CMA-ES", Symposium on Evolutionary Algorithms: Back to the Future and Beyond—Traversing the Ever-Evolving Landscape of Evolutionary Algorithms, Delft, The Netherlands, September 2019
- Nikolaus Hansen: Keynote lecture at the Workshop on Machine-Learning-Assisted Image Formation entitled "From Gradient-Based to Evolutionary Optimization", Nice, France, June 2019
- Anne Auger: Convergence Results of Adaptive Evolution Strategies: an Overview, Lifeware seminar, January 2019
- Anne Auger: COMO-CMA-ES a linearly convergent multi-objective solver, ICCOP conference, Berlin

10.1.6. Leadership within the Scientific Community

- Anne Auger, Elected Member of the ACM-SIGEVO executive board
- Dimo Brockhoff: member of the International Advisory Committee for EMO-2019 in East Lansing, USA
- Anne Auger, member of the conseil de laboratoire of the CMAP, Ecole Polytechnique.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Master: Anne Auger, "Optimization without gradients", 22.5h ETD, niveau M2 (Optimization Master of Paris-Saclay)

Master: Dimo Brockhoff, "Algorithms and Complexity", 36h ETD, niveau M1/M2 (joint MSc with ESSEC "Data Sciences & Business Analytics"), CentraleSupelec, France

Master: Anne Auger and Dimo Brockhoff, "Introduction to Optimization", 31.5h ETD, niveau M2 (MSc Informatique - Parcours Apprentissage, Information et Contenu (AIC)), U. Paris-Saclay, France

Master: Anne Auger and Dimo Brockhoff, "Advanced Optimization", 31.5h ETD, niveau M2 (MSc Informatique - Parcours Apprentissage, Information et Contenu (AIC)), U. Paris-Saclay, France

10.2.2. Tutorials

- Dimo Brockhoff: tutorial on Evolutionary Multiobjective Optimization, 3h ETD, niveau PhD, ACM-GECCO conference, Czech Republic
- Nikolaus Hansen: tutorial "A Practical Guide to Experimentation", 3h ETD, niveau PhD, ACM-GECCO conference, Czech Republic
- Nikolaus Hansen: tutorial on "CMA-ES and Advanced Adaptation Mechanisms", 3h ETD, niveau PhD, ACM-GECCO conference, Czech Republic, with Youhei Akimoto

10.2.3. Supervision

- PhD in progress: Konstantinos Varelas, "Large-Scale Optimization, CMA-ES and Radar Applications", Dec. 2017, Anne Auger and Dimo Brockhoff
- PhD in progress: Cheikh Touré, "Linearly Convergent Multi-objective Stochastic Optimizers", Dec. 2017, Anne Auger and Dimo Brockhoff
- PhD in progress: Paul Dufossé, "Constrained Optimization and Radar Applications", Oct. 2018, Nikolaus Hansen
- PhD in progress: Marie-Ange Dahito, "Mixed-Integer Blackbox Optimization for Multiobjective Problems in the Automotive Industry", Jan 2019, Dimo Brockhoff and Nikolaus Hansen
- PhD in progress: Eugénie Marescaux, Theoretical Analysis of convergence of multi-objective solvers (2019–), supervisor A. Auger

- PhD in progress: Alann Cheral, "Black-box optimization for the optimization of hyperspectral bandwidth for anomaly detection" (2019–), supervisor A. Auger
- Jingyun Yang, Ecole Polytechnique, since September 2019
- Eugenie Marescaux, ENSTA, April–September 2019
- Julien Bonneville, U. Versailles, March–August 2019

10.2.4. Juries

- Anne Auger: member of the **PGMO PhD award** scientific committee. Representing the committee for the PGMO PhD award ceremony at the PGMO days
- Dimo Brockhoff: jury member for the PhD thesis of David Gaudrie, École des Mines de Saint-Étienne, October 2019

10.3. Popularization

- Scientific mediation by Cheikh Touré: presentation on poker in the context of the visit of college and high school students to Inria Saclay

10.3.1. Internal or external Inria responsibilities

- Anne Auger: member of the BCEP of Saclay.
- Dimo Brockhoff: member of the CDT committee of Saclay (since February 2019).

11. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journals

[1] Y. AKIMOTO, N. HANSEN. *Diagonal Acceleration for Covariance Matrix Adaptation Evolution Strategies*, in "Evolutionary Computation", 2019, https://arxiv.org/abs/1905.05885 , forthcoming [*DOI :* 10.1162/EVCO_A_00260], https://hal.inria.fr/hal-01995373

International Conferences with Proceedings

[2] D. BROCKHOFF, N. HANSEN. *The Impact of Sample Volume in Random Search on the bbob Test Suite*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019 [*DOI :* 10.1145/3319619.3326894], https://hal.inria.fr/hal-02171213

[3] D. BROCKHOFF, T. TUŠAR. *Benchmarking Algorithms from the platypus Framework on the Biobjective bbob-biobj Testbed*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019, vol. 7 [*DOI :* 10.1145/3319619.3326896], https://hal.inria.fr/hal-02171136

[4] P. DUFOSSÉ, C. TOURÉ. *Benchmarking MO-CMA-ES and COMO-CMA-ES on the Bi-objective bbob-biobj Testbed*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019 [*DOI :* 10.1145/3319619.3326892], https://hal.inria.fr/hal-02161252

[5] *Best Paper*
N. HANSEN. *A Global Surrogate Assisted CMA-ES*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, ACM, 2019, pp. 664-672 [*DOI :* 10.1145/3321707.3321842], https://hal.inria.fr/hal-02143961.

[6] *Best Paper*

C. TOURÉ, A. AUGER, D. BROCKHOFF, N. HANSEN. *On Bi-Objective convex-quadratic problems*, in "10th International Conference on Evolutionary Multi-Criterion Optimization", East Lansing, Michigan, United States, March 2019, https://arxiv.org/abs/1812.00289 , https://hal.inria.fr/hal-01942159.

[7] C. TOURÉ, N. HANSEN, A. AUGER, D. BROCKHOFF. *Uncrowded Hypervolume Improvement: COMO-CMA-ES and the Sofomore framework*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019, Part of this research has been conducted in the context of a research collaboration between Storengy and Inria [*DOI :* 10.1145/3321707.3321852], https://hal.inria.fr/hal-02103694

[8] T. TUŠAR, D. BROCKHOFF, N. HANSEN. *Mixed-Integer Benchmark Problems for Single-and Bi-Objective Optimization*, in "GECCO 2019 -The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019, submitted to GECCO 2019, https://hal.inria.fr/hal-02067932

[9] K. VARELAS, M.-A. DAHITO. *Benchmarking Multivariate Solvers of SciPy on the Noiseless Testbed*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019 [*DOI :* 10.1145/3319619.3326891], https://hal.inria.fr/hal-02160099

[10] K. VARELAS. *Benchmarking Large Scale Variants of CMA-ES and L-BFGS-B on the bbob-largescale Testbed*, in "GECCO 2019 - The Genetic and Evolutionary Computation Conference", Prague, Czech Republic, July 2019 [*DOI :* 10.1145/3319619.3326893], https://hal.inria.fr/hal-02160106

Other Publications

[11] D. BROCKHOFF, T. TUŠAR, A. AUGER, N. HANSEN. *Using Well-Understood Single-Objective Functions in Multiobjective Black-Box Optimization Test Suites*, January 2019, https://arxiv.org/abs/1604.00359 - ArXiv e-prints, arXiv:1604.00359, https://hal.inria.fr/hal-01296987

[12] O. A. ELHARA, K. VARELAS, D. H. NGUYEN, T. TUŠAR, D. BROCKHOFF, N. HANSEN, A. AUGER. *COCO: The Large Scale Black-Box Optimization Benchmarking (bbob-largescale) Test Suite*, March 2019, https://arxiv.org/abs/1903.06396 - working paper or preprint, https://hal.inria.fr/hal-02068407

[13] S. MAHÉVAS, V. PICHENY, P. LAMBERT, N. DUMOULIN, L. ROUAN, C. SOULIÉ, D. BROCKHOFF, S. LEHUTA, R. LE RICHE, R. FAIVRE, H. DROUINEAU. *A practical guide for conducting calibration and decision-making optimisation with complex ecological models*, December 2019, pdf available at https://www.preprints.org/manuscript/201912.0249/v1 [*DOI :* 10.20944/PREPRINTS201912.0249.V1], https://hal.inria.fr/hal-02418667

References in notes

[14] Y. AKIMOTO, N. HANSEN. *Online model selection for restricted covariance matrix adaptation*, in "International Conference on Parallel Problem Solving from Nature", Springer, 2016, pp. 3–13

[15] Y. AKIMOTO, 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

[16] D. V. ARNOLD, J. PORTER. *Towards au Augmented Lagrangian Constraint Handling Approach for the $(1+1)$-ES*, in "Genetic and Evolutionary Computation Conference", ACM Press, 2015, pp. 249-256

[17] A. ATAMNA, A. AUGER, N. HANSEN. *Linearly Convergent Evolution Strategies via Augmented Lagrangian Constraint Handling*, in "Foundation of Genetic Algorithms (FOGA)", 2017

[18] A. AUGER, N. HANSEN. *Linear Convergence of Comparison-based Step-size Adaptive Randomized Search via Stability of Markov Chains*, in "SIAM Journal on Optimization", 2016, vol. 26, n^o 3, pp. 1589-1624

[19] J. BERGSTRA, R. BARDENET, Y. BENGIO, B. KÉGL. *Algorithms for Hyper-Parameter Optimization*, in "Neural Information Processing Systems (NIPS 2011)", 2011, https://hal.inria.fr/hal-00642998/file/draft1.pdf

[20] V. S. BORKAR. *Stochastic approximation: a dynamical systems viewpoint*, 2008, Cambridge University Press

[21] V. BORKAR, S. MEYN. *The O.D.E. Method for Convergence of Stochastic Approximation and Reinforcement Learning*, in "SIAM Journal on Control and Optimization", January 2000, vol. 38, n^o 2

[22] 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

[23] G. COLLANGE, S. REYNAUD, N. HANSEN. *Covariance Matrix Adaptation Evolution Strategy for Multidisciplinary Optimization of Expendable Launcher Families*, in "13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Proceedings", 2010

[24] J. E. DENNIS, R. B. SCHNABEL. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1983

[25] N. HANSEN, 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

[26] N. HANSEN, A. OSTERMEIER. *Completely Derandomized Self-Adaptation in Evolution Strategies*, in "Evolutionary Computation", 2001, vol. 9, n^o 2, pp. 159–195

[27] J. N. HOOKER. *Testing heuristics: We have it all wrong*, in "Journal of heuristics", 1995, vol. 1, n^o 1, pp. 33–42

[28] F. HUTTER, H. HOOS, K. LEYTON-BROWN. *An Evaluation of Sequential Model-based Optimization for Expensive Blackbox Functions*, in "GECCO (Companion) 2013", ACM, 2013, pp. 1209–1216

[29] 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", 2002, vol. 59, pp. 215–250

[30] D. R. JONES, M. SCHONLAU, W. J. WELCH. *Efficient global optimization of expensive black-box functions*, in "Journal of Global optimization", 1998, vol. 13, n^o 4, pp. 455–492

[31] I. KRIEST, V. SAUERLAND, S. KHATIWALA, A. SRIVASTAV, A. OSCHLIES. *Calibrating a global three-dimensional biogeochemical ocean model (MOPS-1.0)*, in "Geoscientific Model Development", 2017, vol. 10, n⁰ 1, 127 p.

[32] H. J. KUSHNER, G. YIN. *Stochastic approximation and recursive algorithms and applications*, Applications of mathematics, Springer, New York, 2003, http://opac.inria.fr/record=b1099801

[33] P. MACALPINE, S. BARRETT, D. URIELI, V. VU, 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)", July 2012

[34] S. MEYN, R. TWEEDIE. *Markov Chains and Stochastic Stability*, Springer-Verlag, New York, 1993

[35] Y. OLLIVIER, L. ARNOLD, A. AUGER, N. HANSEN. *Information-geometric optimization algorithms: A unifying picture via invariance principles*, in "Journal Of Machine Learning Research", 2016, accepted

[36] T. SALIMANS, J. HO, X. CHEN, I. SUTSKEVER. *Evolution strategies as a scalable alternative to reinforcement learning*, in "arXiv preprint arXiv:1703.03864", 2017

[37] J. SNOEK, H. LAROCHELLE, R. P. ADAMS. *Practical bayesian optimization of machine learning algorithms*, in "Neural Information Processing Systems (NIPS 2012)", 2012, pp. 2951–2959

[38] J. UHLENDORF, A. MIERMONT, T. DELAVEAU, G. CHARVIN, F. FAGES, S. BOTTANI, G. BATT, 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", 2012, vol. 109, n⁰ 35, pp. 14271–14276