Optimizing concentrated solar power plants using adaptive sampling-based derivative-free algorithms

EMMANOUIL KARANTOUMANIS and NIKOLAOS PLOSKAS, University of Western Macedonia, Greece

The optimization of complex real-world systems often presents a challenge when explicit derivatives of the objective function are unavailable. In this paper, we address a real-world black-box optimization problem arising from the design and operation of Concentrated Solar Power (CSP) systems. CSP systems present a unique challenge due to their nonlinear, multi-modal nature, which complicates optimization using traditional gradient-based methods. Derivative-free optimization (DFO) techniques are well-suited to tackle such black-box problems, but even these methods can become impractical when the number of function evaluations required is too large since the evaluation can be expensive. To overcome this limitation, we utilize an adaptive sampling DFO approach that requires a smaller number of function evaluations by intelligently selecting informative points based on surrogate models. The surrogate models are then optimized using derivative-based optimization algorithms to find new sampling points that may be (near-) optimal. We apply this methodology to optimize a CSP problem from the SOLAR benchmark tool, specifically focusing on minimizing the cost of thermal storage, and the total investment cost. The results demonstrate that our adaptive sampling method, ADASNOBFIT, outperforms the well-known SNOBFIT algorithm in terms of solution quality.

CCS Concepts: • **Mathematics of computing → Continuous optimization**; • **Computing methodologies → Continuous simulation**; • **Applied computing → Decision analysis**.

Additional Key Words and Phrases: Black-box optimization, Derivative-free optimization, Surrogate modeling, Adaptive sampling, Concentrated solar power

ACM Reference Format:
Emmanouil Karantoumanis and Nikolaos Ploskas. 2024. Optimizing concentrated solar power plants using adaptive sampling-based derivative-free algorithms. In *Proceedings of 28th Pan-Hellenic Conference on Progress in Computing Informatics with International Participation (PCI 2024).* ACM, New York, NY, USA, 9 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

In real-world optimization problems, it is often the case that the underlying function to be optimized is not explicitly known or analytically tractable [21]. These problems are referred to as "black-box" problems because the internal workings of the system remain hidden, and only input-output relations can be observed. Black-box optimization (BBO) is particularly challenging because traditional optimization methods, such as gradient-based algorithms, rely on explicit derivatives of the objective function, which are not available in black-box scenarios. Instead, derivative-free optimization (DFO) techniques must be employed, relying solely on function evaluations to guide the search for optimal solutions.

Authors' Contact Information: Emmanouil Karantoumanis, e.karantoumanis@uowm.gr; Nikolaos Ploskas, nploskas@uowm.gr, University of Western Macedonia, Kozani, Greece.

PCI 2024, December 13–15, 2024, Greece, Athens

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/XXXXXXX.XXXXXXX

Black-box issues occur in many real-world applications where intricate biological processes, simulations, or physics dictate the system being improved. These issues could include noise, large processing costs, or measurement uncertainties, which would make optimization even more difficult. Examples of black-box problems include tuning machine learning models [8, 25], software tuning [15, 23], and calibrating scientific simulations [7, 20].

To further improve the performance of DFO algorithms, surrogate models are often utilized [19, 28]. In such algorithms, a surrogate model is designed to mimic the objective function based on the previous function evaluations. In case the simulations are expensive, the surrogate model can be used to explore the search space. In order to sample the search space, adaptive sampling techniques are being used to guide the search for new promising points.

One particularly relevant and challenging black-box optimization problem is found in the field of renewable energy, specifically in Concentrated Solar Power (CSP) systems [5, 6, 18, 24, 27]. CSP traps and changes energy emitted from sunlight into thermal energy using mirrors or lenses to concentrate sunlight onto a receiver. From the receiver, thermal energy warms a fluid that uses steam turbines or a different type of heat engine to generate energy. One of its greatest advantages is the ability to store the thermal energy in molten salts within the CSP system. It allows continuous power production by these systems even after sunset or on cloudy days. Additional benefits of this capability include increasing dispatchability in the system to be a stable and reliable source of power to supply load variation, as well as stability in a grid. CSP can also be hybridized to other energy sources, such as the natural gas or biomass-based systems, to improve flexibility and efficiency. This versatility makes CSP important in the transition to renewable energy and the decarbonization of electricity grids.

However, because of the intricate nature of their components and interconnections, CSP systems provide challenging black-box issues for both design and operational optimization [3, 17]. The design and arrangement of solar collectors, the effectiveness of heat transfer systems, the capacity and control of thermal storage, and the power cycle's overall efficiency are just a few of the interrelated elements that affect a CSP system's performance. System performance is further complicated by environmental factors as wind speed, ambient temperature, and sun irradiation. These factors are typically modeled through detailed physics-based or empirical simulations that do not provide explicit gradient information, which presents a significant challenge for optimization. As a result, traditional gradient-based optimization methods struggle with the inherent non-linearities, noise, and uncertainties of these simulations, particularly due to the lack of differentiable models. CSP systems often exhibit discontinuities in their performance landscape, such as phase changes in heat transfer fluids or varying efficiency thresholds in power cycles, making gradient-based methods unsuitable for robust optimization.

Given these complexities, DFO methods are highly suitable for addressing these challenges. DFO techniques, which do not rely on gradient information, can explore the design space efficiently even when the performance function is noisy, discontinuous, or involves high-dimensional variables. These methods are particularly effective in scenarios where evaluations are expensive or when model fidelity makes computation prohibitive, allowing for a more flexible and adaptive approach to optimizing CSP systems. By leveraging DFO, it becomes possible to identify (near-) optimal configurations and operational strategies that maximize energy output, efficiency, and cost-effectiveness, while accounting for the uncertainty and variability inherent in CSP operations.

Even conventional DFO approaches frequently need a significant number of function evaluations to converge to a satisfactory solution, even though they are appropriate for black-box problems. When each function evaluation is computationally costly, as is the case in many real-world applications, such as CSP system simulations, this is especially problematic. These methods' practical utility is limited by their high computing cost, particularly in situations where time or resources are limited. In order to overcome this difficulty, we employ a DFO algorithm with warm start points and an adaptive sampling technique. By carefully choosing the most instructive locations for evaluation and using surrogate models to approximate the target function, this method seeks to reduce the number of function evaluations. By refining the surrogate model iteratively based on new data, the algorithm can guide the search toward promising regions of the solution space, achieving high-quality results in a significantly smaller number of iterations.

In this paper, we present the results of utilizing an adaptive sampling DFO algorithm to optimize a CSP system. The problem is approached as a black-box one, and our goal is to demonstrate the effectiveness of the adaptive sampling DFO algorithm compared to a classic DFO algorithm in finding near-optimal solutions for a highly nonlinear, multi-modal CSP problem while minimizing the number of expensive function evaluations.

The remainder of the paper is organized as follows. An overview of the related work is presented in Section 2. In Section 3, we describe the methodology employed in this study. Section 4 provides the obtained results and insights. Finally, conclusions are provided in Section 5.

2 Related work

In many domains, optimization is essential, especially when it comes to managing intricate systems with functional boundaries and constraints. It can be difficult to simulate production processes effectively in many industrial applications. Consequently, there is an increasing demand for BBO algorithms that can solve real-world issues without the requirement for exact production models.

Due to its many difficulties, CSP systems are well suited for black-box optimization, particularly when costly function evaluations are required. Complex characteristics like as receiver design, heat storage configurations, and heliostat field patterns must frequently be optimized for these systems. With the help of thermal energy storage devices, CSP technology uses mirrors or lenses to focus sunlight, producing heat that is then transformed into electricity. These plants must manage variables such as energy storage, solar field operation, and turbine performance under varying solar radiation conditions. Because performance evaluation of every design is computationally expensive since detailed simulations must be run for different environmental and operational conditions, nonlinear interactions between the design variables and performance metrics combined with the high cost of the simulations make the CSP optimization problems hard to solve with traditional methods. This has motivated the use of black-box approaches. Other works have used black-box optimization algorithms for optimizing CSP plants, which will be reviewed in this section.

A local search technique called rqlif was created by Manno et al. [17] to optimize expensive black-box issues, with a special emphasis on the start-up stage of CSP plants. To improve solution efficiency, this derivative-free optimization approach uses linear implicit filtering and a regularized quadratic model. The numerical experiments conducted by the experimenters revealed that rqlif was successful in solving 76.63

Hamilton et al. [9] developed a methodology for optimizing the design of CSP and photovoltaic (PV) hybrid systems using NLopt's DFO algorithms [12]. Their approach evaluates the financial feasibility and performance of these systems under various time-of-delivery pricing structures. The study found that the optimized designs could improve the base case power purchase agreement price by 15% to 21% while achieving capacity factors between 50% and 62%. These results highlight the potential for enhanced efficiency and cost-effectiveness in renewable energy solutions.

Luo et al. [16] developed a robust optimization framework to handle uncertainties affecting the design of a molten salt solar power tower plant. The study used a combination of Monte Carlo simulation for uncertainty propagation and a simulated annealing algorithm to solve the multi-objective optimization problem. The objectives include minimizing both the expected value and the standard deviation of the levelized cost of energy. The results reveal a trade-off between minimizing economic cost and reducing risk under uncertainty, with a final optimal solution yielding a 23.09 c/kWh expected levelized cost of energy and 1.25 c/kWh standard deviation, reducing economic risk by 17.22% compared to the deterministic design. The work also employs Sobol's global sensitivity analysis to identify critical parameters like direct solar radiation and heliostat field costs as major influencers on the model output, which helps guide the optimal design.

Cox et al. [3] utilize formal design-of-experiment sampling designs and a Bayesian optimization algorithm for optimizing utility-scale solar plants with energy storage in order to maximize economic performance. The study combines the National Renewable Energy Laboratory's (NREL) System Advisor Model (SAM) [2] for simulating plant operations with a revenue-maximizing mixed-integer linear program for dispatch optimization. A Bayesian optimization algorithm is used to design plants with photovoltaics, CSO, and energy storage. Results show improvements of 6–19% in lifetime benefit-to-cost ratios through optimal system sizing.

The SOLAR tool [1] is designed as a black-box simulator to benchmark optimization algorithms for CSP systems. It simulates real-world scenarios, creating a set of benchmark problems based on different operational configurations and constraints of the power plant. These benchmark problems are designed to test optimization algorithms on both continuous and mixed-integer variables, challenging them to maximize energy efficiency and minimize operational costs while respecting physical constraints like temperature limits and material capacities. The benchmark problems include tasks such as optimizing the turbine's operational schedule or managing the heat storage system. Andrés-Thió et al. [1] employed GA, CMA-ES [10], and NOMAD [14] algorithms to find near-optimal solutions. The results showed that NOMAD effectively navigated the feasible solution space, often achieving high-quality solutions with fewer function evaluations; however, these algorithms required thousands of function evaluations.

In our previous work [13], we introduced an adaptive sampling procedure utilizing the SNOBFIT [11] algorithm. This approach was evaluated using a comprehensive benchmark consisting of 776 continuous optimization problems, each characterized by bound constraints. Our extensive testing demonstrated the efficacy of this method, successfully solving 93% of the problems. In this paper, we apply the adaptive sampling procedure to solve a real-world CSP problem with a few function evaluations.

3 Methodology

Our approach addresses black-box optimization using an adaptive sampling method, breaking the process into stages, as shown in Figure 1. As n indicates the problem's dimension, the process entails producing points and building surrogate models until a termination condition based on the number of points (n_{final}) or the R^2 score threshold is satisfied.

The method starts with an initial design of experiments by Latin Hypercube Sampling (LHS) to generate $2n$ points in the search space. LHS provides a diverse and well-scattered set of initial points. This reduces the possibility of local optima and gives a good overview of the optimization landscape. Next, we construct a surrogate model to approximate the black-box function, which is refined iteratively using new points generated by the DFO solver SNOBFIT. This surrogate model helps search space exploration and guides the optimization to promising areas of minimum points. For surrogate modeling, we utilize the tool ALAMO [4], which mainly employs methods such as LASSO regularization while building the model. Moreover, local and global Derivative-based optimization (DBO) solvers are employed for optimization. The use of BARON [22] guarantees the investigation of the global minimum while refining the solution with the use of IPOPT [26] as a local optimizer.

The adaptive sampling procedure employs a DFO approach, allowing exploration of the search space without relying on derivatives. By integrating an Error Maximization Strategy (EMS) 1, the

Fig. 1. Adaptive sampling methodology

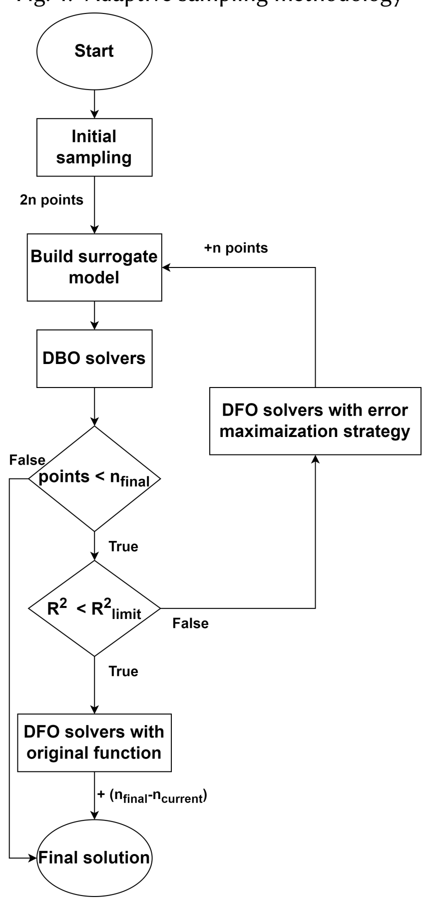

algorithm targets areas with the largest discrepancies between the surrogate model's predictions and the actual outputs, thereby enhancing efficiency and reducing function evaluations. Each iteration generates new points based on the EMS, refining the surrogate model and establishing a feedback loop that improves its predictive accuracy. If the surrogate model's R^2 score drops below a threshold, indicating inadequate accuracy, the algorithm shifts to directly evaluating the original black-box function, maximizing the use of the remaining function calls within the budget.

$$\max_{x} \left(\frac{y(x) - \widehat{y}(x)}{y(x)} \right)^2 \tag{1}$$

where:

- $y(x)$ represents the actual output value of the black-box problem
- $\widehat{y}(x)$ represents the surrogate model's output value

4 Experimental results

The SOLAR tool provides 10 benchmark black-box problems for optimization, each designed to challenge optimization algorithms in various contexts, including both constrained and unconstrained

Problem	Total variables	Continuous	Integer	Fix
3	11	9	2	x1=4.6, x3=71.9, x5=6.2, x7=38.3, x8=2.3, x9=11.2, x16=275, x18=0.0165, x19=0.018, x20=0.019
4	17	11	6	x1=6, x3=70, x5=8, x7=45, x8=2.5, x9=6, x16=520, x18=0.0165, x19=0.018, x20=0.017, x22=0.0155, x23=0.016
10	5	5	-	-

Table 1. Variable analysis

scenarios. Out of these 10 problems, the problem we focus on is the 10th, as it is the only fully unconstrained problem. Problem 10 specifically deals with minimizing the cost of storage, a critical factor in optimizing the performance and economic feasibility of renewable energy systems, such as CSP.

This problem is an adaptation of problem 6, which also focuses on minimizing storage costs but operates within a constrained environment. While problem 10 does not set explicit limits on the variables of the problem system, it uses a penalty mechanism to allow violations of the constraints as an implicit part of the optimization objective. This provides a smoother, more flexible search space where the optimization process may investigate areas that in conventional constraint-based formulations would have been considered infeasible but at the cost of higher penalties. Therefore, the algorithm will explore and exploit effectively by trading off storage cost minimization with violations in constraints.

To expand the experimental study, we modified two additional SOLAR tool problems, specifically problems 3 and 4. The objective for problems 3 and 4 is to minimize the total investment cost while satisfying the demand and respecting a maximum field size. We selected these problems due to their focus on cost minimization with mixed-integer variables. As with the previous case, we applied penalty techniques to non-valid constraints. Some variables in these problems were part of hidden constraints that could not be violated. To address this, we converted these variables into constants, assigning them values from the best solution identified by NOMAD.

The problems involve continuous or mixed-integer variables representing various design parameters of a CSP system. Some of these variables include the central receiver outlet temperature, affecting thermal energy available for storage and power generation; the height and diameter of the hot storage tank, influencing energy capacity; the insulation thickness of the hot, and cold storage tanks which is crucial for minimizing heat loss and maintaining system efficiency, etc. In Table 1 we present the total number of variables for each problem, detailing the number of continuous and integer variables, as well as the constant values assigned to specific variables that are subject to hidden constraints.

We compared the performance of our suggested approach with a well-known DFO method in order to fully evaluate its effectiveness. In particular, we conducted five iterations of each experiment for each algorithm to compare the suggested approach, ADASNOBFIT, to SNOBFIT. The average values from these five runs are shown in the results below, guaranteeing solid and statistically significant results that demonstrate the efficiency and dependability of our adaptive sampling strategy in comparison to stand-alone DFO techniques.

We applied a uniform termination criterion of $n_{final} = 12 \times n$ iterations, where n represents the problem's dimensionality. For the benchmark problem with a dimension of 5, this resulted in

Problem	SNOBFIT optimal value	ADASNOBFIT optimal value	SNOBFIT time	ADASNOBFIT time
3	4365.77	975.23	710	889
4	4,779.16	4,699.47	1,202	1,752
10	87.23	53.01	3,606	4,050

Table 2. Results

$n_{final} = 5 \times 12 = 60$ function evaluations, ensuring both algorithms had equal opportunities to explore the optimization landscape.

As previously mentioned, the adaptive sampling process continues until the number of sampled points exceeds n_{final} or the R^2 score criteria are met. The R^2 score must satisfy two conditions: $R^2 > 0.65$ (indicating the model's performance has degraded) and R^2previous $-$ R^2current < 0.2 (indicating the current surrogate model performs significantly worse than the previous one). If either condition is unmet, the adaptive sampling process stops, and the DFO algorithm proceeds with the original black-box function for the remaining iterations.

SNOBFIT and ADASNOBFIT are clearly compared in Table 2 with respect to execution time and solution quality for problems 3, 4, and 10. In terms of solution accuracy ADASNOBFIT consistently outperforms SNOBFIT, obtaining superior values across all problems, demonstrating its effectiveness in optimizing within strict evaluation bounds. For instance, ADASNOBFIT achieves a value of 53.01 in issue 10, which is significantly closer to the best-known answer of 42.41 (obtained by NOMAD after 2,000 evaluations) than SNOBFIT's 87.23. This demonstrates how ADASNOBFIT can produce near-optimal solutions with fewer evaluations.

On the other hand, SNOBFIT is faster in execution, as shown by the shorter times across all problems. The additional time required by ADASNOBFIT is primarily due to the extra computational cost after the first $2n$ iterations, as building and solving the surrogate model using ALAMO, BARON, and IPOPT introduces an overhead. Despite the additional time, ADASNOBFIT's superior solution quality across all problems suggests a valuable trade-off between computation time and solution accuracy.

Figure 2 compares the performance of the SNOBFIT and ADASNOBFIT algorithms for the 3 problems. The x-axis represents the number of function evaluations, while the y-axis shows the solution value. Initially, both algorithms start with the same initial point however, ADASNOBFIT quickly converges to a much lower solution after the construction of the first surrogate model. The ADASNOBFIT method demonstrates better performance by converging faster and reaching a lower final solution value compared to SNOBFIT.

5 Conclusions

In this paper, we applied an adaptive sampling DFO approach to a challenging black-box optimization problem in CSP system design. The results demonstrate that the proposed ADASNOBFIT method outperforms the SNOBFIT algorithm for each of the three CSP benchmark problems. Notably, in problem 10 the proposed ADASNOBFIT methodology significantly outperformed SNOBFIT, achieving a solution just 10.6 away from the best-known value of 42.41 obtained by the NOMAD algorithm, but with only 60 function evaluations compared to NOMAD's 2,000, demonstrating the efficiency and effectiveness of our approach. Although ADASNOBFIT required more computation time due to surrogate modeling overhead, the trade-off between function evaluation cost and time complexity remains favorable in scenarios where each evaluation is expensive, such as in large-scale CSP simulations. The study shows that adaptive sampling DFO can be a powerful tool

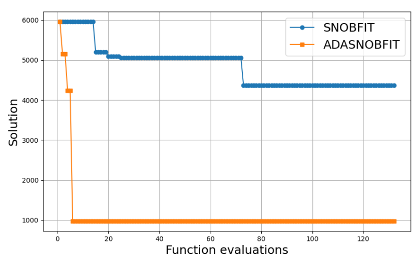

(a) Problem 3

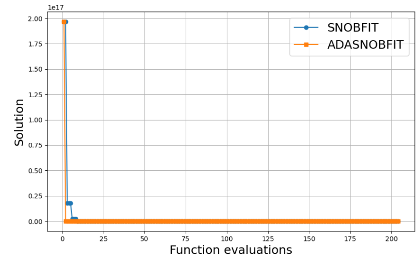

(b) Problem 4

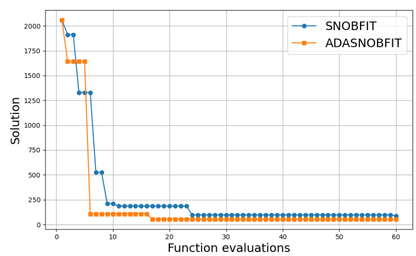

(c) Problem 10

Fig. 2. Solution comparison of SNOBFIT and ADASNOBFIT

for optimizing CSP systems, especially in real-world applications where the evaluation budget is limited. Future work will focus on extending this approach to more complex and realistic CSP problems, including multi-objective optimization and the incorporation of additional operational constraints.

Acknowledgment This work has been co-financed by the National Strategic Reference Framework 2021-2027 of the European Union and Greek national funds through the Operational Program Competitiveness (project code MIS 6001586).

References

[1] N. Andrés-Thió, C. Audet, M. Diago, A.E. Gheribi, S. Le Digabel, X. Lebeuf, M. Lemyre Garneau, and C. Tribes. 2024. `solar`: *A solar thermal power plant simulator for blackbox optimization benchmarking*. Technical Report G-2024-37. Les cahiers du GERAD. http://www.optimization-online.org/XXX

[2] Nate Blair, Aron P Dobos, Janine Freeman, Ty Neises, Michael Wagner, Tom Ferguson, Paul Gilman, and Steven Janzou. 2014. *System advisor model, sam 2014.1. 14: General description*. Technical Report. National Renewable Energy Lab.(NREL), Golden, CO (United States).

[3] John Cox, William T Hamilton, Alexandra M Newman, and Janna Martinek. 2023. Optimal sizing and dispatch of solar power with storage. *Optimization and Engineering* 24, 4 (2023), 2579–2617.

[4] Alison Cozad, Nikolaos V Sahinidis, and David C Miller. 2014. Learning surrogate models for simulation-based optimization. *AIChE Journal* 60, 6 (2014), 2211–2227.

[5] Ershun Du, Ning Zhang, Bri-Mathias Hodge, Qin Wang, Chongqing Kang, Benjamin Kroposki, and Qing Xia. 2018. The role of concentrating solar power toward high renewable energy penetrated power systems. *IEEE Transactions on Power Systems* 33, 6 (2018), 6630–6641.

[6] Tobias Fichter, Rafael Soria, Alexandre Szklo, Roberto Schaeffer, and Andre FP Lucena. 2017. Assessing the potential role of concentrated solar power (CSP) for the northeast power system of Brazil using a detailed power system model. *Energy* 121 (2017), 695–715.

[7] Alexander IJ Forrester. 2010. Black-box calibration for complex-system simulation. *Philosophical Transactions of The Royal Society A: Mathematical, Physical and Engineering Sciences* 368, 1924 (2010), 3567–3579.

[8] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and David Sculley. 2017. Google vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1487–1495.

[9] William T Hamilton, Michael J Wagner, Alexandra M Newman, and Robert J Braun. 2020. Black-box optimization for design of concentrating solar power and photovoltaic hybrid systems with optimal dispatch decisions. In *AIP Conference Proceedings*, Vol. 2303. AIP Publishing.

[10] Nikolaus Hansen. 2016. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).

[11] Waltraud Huyer and Arnold Neumaier. 2008. SNOBFIT–stable noisy optimization by branch and fit. *ACM Trans. Math. Software* 35, 2 (2008), 1–25.

[12] Steven G. Johnson. 2007. The NLopt nonlinear-optimization package. https://github.com/stevengj/nlopt.

[13] Emmanouil Karantoumanis and Nikolaos Ploskas. 2024. Improving derivative-free optimization algorithms through an adaptive sampling procedure. *Results in Control and Optimization* 16 (2024), 100460.

[14] Sébastien Le Digabel. 2011. Algorithm 909: NOMAD: Nonlinear Optimization with the MADS Algorithm. *ACM Trans. Math. Softw.* 37, 4, Article 44 (Feb. 2011), 15 pages. https://doi.org/10.1145/1916461.1916468

[15] Jianfeng Liu, Nikolaos Ploskas, and Nikolaos V Sahinidis. 2019. Tuning BARON using derivative-free optimization algorithms. *Journal of Global Optimization* 74 (2019), 611–637.

[16] Yan Luo, Zhiyuan Wang, Jiamin Zhu, Tao Lu, Gang Xiao, Fengming Chu, and Ruixing Wang. 2022. Multi-objective robust optimization of a solar power tower plant under uncertainty. *Energy* 238 (2022), 121716.

[17] Andrea Manno, Edoardo Amaldi, Francesco Casella, and Emanuele Martelli. 2020. A local search method for costly black-box problems and its application to CSP plant start-up optimization refinement. *Optimization and Engineering* 21 (2020), 1563–1598.

[18] R Mena, Rodrigo Escobar, Á Lorca, M Negrete-Pincetic, and DJAE Olivares. 2019. The impact of concentrated solar power in electric power systems: A Chilean case study. *Applied Energy* 235 (2019), 258–283.

[19] Juliane Müller, Christine A Shoemaker, and Robert Piché. 2013. SO-MI: A surrogate model algorithm for computationally expensive nonlinear mixed-integer black-box global optimization problems. *Computers & Operations Research* 40, 5 (2013), 1383–1400.

[20] Amir Parnianifard, A Azfanizam, MKAM Ariffin, M Ismail, and N Ebrahim. 2019. Recent developments in metamodel based robust black-box simulation optimization: An overview. *Decision Science Letters* 8, 1 (2019), 17–44.

[21] Nikolaos Ploskas and Nikolaos V Sahinidis. 2022. Review and comparison of algorithms and software for mixed-integer derivative-free optimization. *Journal of Global Optimization* 82 (2022), 433–462.

[22] Nikolaos V Sahinidis. 1996. BARON: A general purpose global optimization software package. *Journal of Global Optimization* 8, 2 (1996), 201–205.

[23] Benjamin Sauk, Nikolaos Ploskas, and Nikolaos Sahinidis. 2020. GPU parameter tuning for tall and skinny dense linear least squares problems. *Optimization Methods and Software* 35, 3 (2020), 638–660.

[24] Rafael Soria, André FP Lucena, Jan Tomaschek, Tobias Fichter, Thomas Haasz, Alexandre Szklo, Roberto Schaeffer, Pedro Rochedo, Ulrich Fahl, and Jürgen Kern. 2016. Modelling concentrated solar power (CSP) in the Brazilian energy system: A soft-linked model coupling approach. *Energy* 116 (2016), 265–280.

[25] Ryan Turner, David Eriksson, Michael McCourt, Juha Kiili, Eero Laaksonen, Zhen Xu, and Isabelle Guyon. 2021. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. In *NeurIPS 2020 Competition and Demonstration Track*. PMLR, 3–26.

[26] Andreas Wächter and Lorenz T Biegler. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106 (2006), 25–57.

[27] Lingxiang Yao, Yang Wang, and Xianyong Xiao. 2023. Concentrated solar power plant modeling for power system studies. *IEEE Transactions on Power Systems* (2023).

[28] Jianyuan Zhai and Fani Boukouvala. 2023. Surrogate-based branch-and-bound algorithms for simulation-based black-box optimization. *Optimization and Engineering* 24, 3 (2023), 1463–1491.

Received 20 October 2024