

Escaping local minima with derivative-free methods: a numerical investigation

Coralia Cartis* Lindon Roberts* Oliver Sheridan-Methven*

28th December 2018

Abstract

We apply a state-of-the-art, local derivative-free solver, Py-BOBYQA [6], to global optimization problems, and propose an algorithmic improvement that is beneficial in this context. Our numerical findings are illustrated on a commonly-used test set of global optimization problems [1] and associated noisy variants, and on hyperparameter tuning for the machine learning test set MNIST [40]. As Py-BOBYQA is a model-based trust-region method, we compare mostly (but not exclusively) with other global optimization methods for which (global) models are important, such as Bayesian optimization and response surface methods; we also consider state-of-the-art representative deterministic and stochastic codes, such as DIRECT and CMA-ES. We find Py-BOBYQA to be competitive with global solvers that are provably designed for finding global optima, for all accuracy/budget regimes, in both smooth and noisy settings. In particular, Py-BOBYQA variants are best performing for smooth and multiplicative noise problems in high-accuracy regimes. As a by-product, some preliminary conclusions can be drawn on the relative performance of the global solvers we have tested with default settings.

Keywords: derivative-free optimization, global optimization, trust region methods.

Mathematics Subject Classification: 65K05, 90C30, 90C26, 90C56

1 Introduction

Derivative-Free Optimization (DFO) refers to a large class of methods that are designed to find local optima without needing derivative information to be provided or to even exist, a ubiquitous practical situation arising in problems where the often black-box function evaluations may be expensive or inaccurate/noisy [9]. Global Optimization (GO) algorithms, whether derivative-based or derivative-free, tackle the (NP-hard) question of finding global optima, a much more computationally intense effort but a crucial requirement for some applications [27]. Both classes of methods aim to go beyond the remit of the successful and scalable local derivative-based optimization solvers that, by construction, get easily trapped in/close to local optima; this larger applicability comes at the cost of reduced scalability, especially in the case of generic GO algorithms.

Despite only local theoretical guarantees of convergence, DFO methods have long been regarded as a more suitable alternative for GO than local derivative-based methods. Indeed, Rios and Sahinidis [37] recently performed extensive benchmarking of various GO and DFO solvers, and found that in some instances, such as when close to a solution, state-of-the-art DFO methods such as NEWUOA [33] can be comparable, or even more efficient than GO solvers (to reach the same accuracy), but that generally, GO is better than DFO. We have recently improved upon the performance of NEWUOA, with our solver Py-BOBYQA [6], which is a Python-conversion

*Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter, Woodstock Road, Oxford, OX2 6GG, United Kingdom (`cartis`, `roberts1`, `oliver.sheridan-methven@maths.ox.ac.uk`). This work was supported by the EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling (EP/L015803/1) in collaboration with the Numerical Algorithms Group Ltd.

and extension of BOBYQA¹[34]. Our improvements to BOBYQA that are relevant here take the form of *restarts*; we have found in [6] that they significantly improve the performance of BOBYQA in the presence of noise, and also sometimes give improved accuracy in the noiseless case on standard nonlinear least-squares problems. Encouraged by this, in this paper, we further modify the restart procedure in a way that is (even more) advantageous to global optimization, and test both the usual Py-BOBYQA (with restarts) and its GO modification against state of the art GO codes—on standard GO test problems and on a parameter tuning problem from machine learning—with encouraging results, as we describe below.

Literature review There are several types of DFO algorithms: pattern or direct search [25], the well-known simplex/Nelder-Mead algorithm [42], linesearch methods [32], implicit filtering [24], conjugate directions [32], model-based trust region methods [9] and many more, including extensions to constrained problems; Chapter 9 in [31] gives a good survey and the reference textbook [9] provides the details. It is well-known that model-based trust region methods are competitive within existing DFO methods and solvers [10, 29], and we have recently compared different methods with similar conclusions [38, 6]. Thus, in this work, we focus on model-based trust-region DFO methods for unconstrained and bound-constrained optimization problems, in particular on Py-BOBYQA variants [6], which improve upon BOBYQA [34]. These methods build, at each iteration, a *local quadratic model* of the objective function by interpolation of function values, that is minimized over a trust region ball, to find an improvement on the current, best-available, estimate of the solution. When no improvement can be achieved, the size of the trust region and the geometry of the interpolation set are adjusted in the run of the algorithm. The evaluation cost of these methods is larger than the number of variables for the initial iteration (when the first local quadratic model is set up), but it is only order 1 subsequently, as usually only one, or a small number of, points leaves and enters the interpolation set for the remaining iterations. Py-BOBYQA has additional features such as restarts, that we describe below.

GO methods typically fall in two broad categories: deterministic such as branch-and-bound, Lipschitz optimization or response surface/surrogate methods, and stochastic such as random search, evolutionary methods, simulated annealing and Bayesian methods; there are also mixed variants, with both random and deterministic features (such as multi-start methods, as well as variants of surrogate methods). For our comparisons here, we focus mostly on GO methods that also use simplified models of the objective in their construction. Hence we consider state-of-the-art Bayesian optimization solvers, that use (global) Gaussian or other models, and response surface methods that use radial basis function surrogate models; and also SNOBFIT, that is similar to Py-BOBYQA in using local quadratic models, combined with a branching strategy to enforce global searches. For reference, we also select the well-known CMA-ES [15, 16] and DIRECT [22] methods, as representatives for the evolutionary and deterministic classes of methods, respectively².

In terms of test problems, we have selected a class of 50 standard GO test problems [1], of dimensions up to 50 variables, and we also generate noisy variants of these problems by random perturbations, since the stochastic regime is very much within the design scope of both DFO, Bayesian and response surface methods. This test set is representative and common to usual GO solvers such as DIRECT, but probably less familiar to Bayesian solvers which are very much developed and used in the machine learning community. To play to the strengths of Bayesian solvers, we also choose a case study in machine learning, namely, an algorithm parameter tuning problem for training classifiers on the MNIST image data set.

The benchmarking paper [37] compares a large number of DFO and GO solvers on a large test set, as mentioned above, but do not include Bayesian solvers as extensively as we consider here (but do include the related response surface methods). The results in [13] compare model-based DFO with Bayesian solvers on machine learning test sets, so as to emphasize that there is no need to use GO solvers for parameter tuning of training algorithms since local DFO gives as good low-accuracy solutions as Bayesian solvers, with reduced computational and evaluation cost. Our aim here is to report preliminary testing of our recent and significantly improved model-based

¹Note that NEWUOA and BOBYQA are very similar codes, with the latter allowing bound constraints.

²Due to our use of Python for the DFO code, for the sake of like-for-like comparisons, we were limited in our choice of GO solvers to those written in or interfacing to Python.

DFO solver (Py-BOBYQA) in similar (but reduced size) contexts as these two papers, so as to ascertain its suitability or otherwise for GO problems.

Our contributions and findings Py-BOBYQA has a restart mechanism that whenever the algorithm no longer makes much progress, it resets the (small) size of trust region radius to its initial (large) value. It also regenerates completely or partially the tightly-clustered interpolation set to be more spread out in the enlarged trust region. The restarts procedure was designed with noisy problems in mind, where we often observed that the trust region gets too small too quickly, before the solver reaches desired accuracy. In [6], we found that for local optimization, restarts—and the associated auto-detection mechanism—help improve the accuracy of solutions for noisy problems, and do not adversely affect, and may even help, solver performance on smooth problems; our testing for this last point was limited to nonlinear least squares problems [29]. Here we are proposing an additional improvement, *adaptive restarts*, with global optimization in mind: every time a restart is called, the trust region size is not just reset to its original value used at the start of the algorithm, but possibly increased, to allow a larger local region of search and optimization. We test both the default restart configurations and the adaptive restart modification of Py-BOBYQA on the GO test problems and the machine learning test set. We find that for low accuracy solutions on the GO test sets, most solvers have similar performance on a given budget, for both smooth and noisy problems. Differences start appearing for longer budgets/high accuracy settings; here, adaptive Py-BOBYQA variants are the most efficient solvers for smooth problems on a given evaluation budget, while for noisy problems, an adaptive Py-BOBYQA variant (with full quadratic models) is amongst the more efficient solvers, competitive with DIRECT and the surrogate solver PySOT [11]. For the MNIST test set, due to the high cost of the evaluations, only low accuracy solutions were attempted. In this setting, PySOT and DIRECT are the most efficient solvers, while adaptive Py-BOBYQA (with full quadratic models) performs acceptably. Overall, we find that including the adaptive restart strategy in Py-BOBYQA is advantageous for escaping local minima, and that there is reasonably small variability in performance due to the choice of starting point. In terms of the other solvers, DIRECT and PySOT are most competitive across all problems types.

Structure of the paper Section 2 describes the Py-BOBYQA algorithm and its improvements, Section 3 summarizes the main solvers we compare with, while Section 4 presents the testing methodology, including test problems and measuring performance of solvers. Section 5 presents our numerical findings, comparing the different Py-BOBYQA variants, and then the best performing one with the remaining solvers; we give our conclusions in Section 6.

2 Py-BOBYQA: a local DFO solver with improvements

The Py-BOBYQA algorithm is a model-based derivative-free optimization method, where, as is typical for such methods, a local model for the objective is constructed by interpolation and minimized on each iteration [9]. More specifically, Py-BOBYQA [6] is based on Powell’s BOBYQA [34], which uses underdetermined quadratic interpolation for model construction. At each iteration k , we have a collection of points $Y_k \subset \mathbb{R}^n$ with $n + 1 \leq |Y_k| \leq (n + 1)(n + 2)/2$, and we construct a local model for the objective

$$f(\mathbf{x}_k + \mathbf{s}) \approx m_k(\mathbf{s}) = c_k + \mathbf{g}_k^\top \mathbf{s} + \frac{1}{2} \mathbf{s}^\top H_k \mathbf{s}, \quad (2.1)$$

satisfying the interpolation conditions

$$m_k(\mathbf{y}_t - \mathbf{x}_k) = f(\mathbf{y}_t), \quad \text{for all } \mathbf{y}_t \in Y_k. \quad (2.2)$$

If $|Y_k| < (n + 1)(n + 2)/2$, then there are multiple solutions to (2.2), so we use the remaining degrees of freedom by choosing $H_k = 0$ if $|Y_k| = n + 1$, and solving

$$\min_{c_k, \mathbf{g}_k, H_k} \|H_k - H_{k-1}\|_F^2 \quad \text{subject to (2.2)}, \quad (2.3)$$

otherwise, where $\|\cdot\|_F$ is the matrix Frobenius norm and we have the convention $H_{-1} = 0$. The value of $|Y_k|$ is a user-specified input, which defaults to $2n + 1$ for smooth problems and $(n + 1)(n + 2)/2$ for noisy problems.

Once we have a method for constructing local models (2.1), we calculate new steps using a trust-region method [9], which ensures global convergence (i.e. convergence to a local minimum regardless of starting point). That is, we maintain a parameter $\Delta_k > 0$ and compute a new step by solving

$$\mathbf{s}_k \approx \arg \min_{\|\mathbf{s}\| \leq \Delta_k} m_k(\mathbf{s}). \quad (2.4)$$

Efficient algorithms exist for solving (2.4) approximately (e.g. [8]). We then evaluate the objective $\mathbf{x}_k + \mathbf{s}_k$, and determine if the step produce a sufficient decrease in the objective, by computing:

$$r_k = \frac{\text{actual decrease}}{\text{expected decrease}} := \frac{f(\mathbf{x}_k) - f(\mathbf{x}_k + \mathbf{s}_k)}{m_k(\mathbf{0}) - m_k(\mathbf{s}_k)}. \quad (2.5)$$

If r_k is sufficiently large, we accept the step ($\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{s}_k$) and increase Δ_k , otherwise we reject the step ($\mathbf{x}_{k+1} = \mathbf{x}_k$) and decrease Δ_k .

The last key component of the algorithm is the maintenance of the interpolation set Y_k . We want Y_k to consist of points close to \mathbf{x}_k , with ‘good’ geometry. The quality of Y_k is calculated based on Lagrange polynomials; that is, the polynomial functions $\{\ell_1(\mathbf{y}), \dots, \ell_p(\mathbf{y})\}$ for $p = |Y_k|$ defined by

$$\ell_t(\mathbf{y}_s) = \begin{cases} 1, & t = s, \\ 0, & t \neq s, \end{cases} \quad \forall s, t = 1, \dots, p. \quad (2.6)$$

Given these Lagrange polynomials, the measure of the quality of Y_k is given by the following definition.

Definition 2.1 (Λ -poised; Chapter 6, [9]). For $B \subset \mathbb{R}^n$ and $\Lambda > 0$, the set Y_k with is Λ -poised in B if

$$\max_{t=1, \dots, p} \max_{\mathbf{y} \in B} |\ell_t(\mathbf{y})| \leq \Lambda, \quad (2.7)$$

for all $\mathbf{y} \in B$, where $\{\ell_1(\mathbf{y}), \dots, \ell_p(\mathbf{y})\}$ are the Lagrange polynomials for Y_k .

Updating Y_k proceeds as follows: at each iteration we add $\mathbf{x}_k + \mathbf{s}_k$ to Y_k (and remove a point far from \mathbf{x}_{k+1}), and at certain times we must also replace a point in Y_k with a new point designed solely to improve the Λ -poisedness of Y_k (rather than decreasing the objective).

The complete Py-BOBYQA algorithm The basic Py-BOBYQA algorithm is summarized in Algorithm 1; a complete description can be found in [34, 6]. Its implementation also allows optional bound constraints $\mathbf{a} \leq \mathbf{x} \leq \mathbf{b}$; the main impact of this is the need for a specialized routine for solving (2.4) with bound constraints, which we take from [34]. When bound constraints are supplied, Py-BOBYQA can internally shift and scale the variables to lie in $[0, 1]$.

2.1 Multiple restarts feature

An important feature of Py-BOBYQA, originally designed to improve its robustness for noisy problems [6] is the ability to restart the solver. When the objective is noisy, this stops Py-BOBYQA from stagnating in regions far from local minima; however, we will see that it is also useful for escaping from local minima. A restart is triggered in four scenarios:

1. The trust region radius Δ_k becomes sufficiently small;
2. The decrease in the objective over consecutive iterations is too slow;
3. If when the objective values $f(\mathbf{y})$ for $\mathbf{y} \in Y_k$ are all within some (user-specified) noise level; or
4. Over consecutive iterations, the trust region radius is consistently being decreased, and the changes $\log \|\mathbf{g}_k - \mathbf{g}_{k-1}\|$ and $\log \|H_k - H_{k-1}\|_F$ are consistently increasing.

Algorithm 1 Py-BOBYQA for general minimization; see [34, 6] for full details.

Input: Starting point $\mathbf{x}_0 \in \mathbb{R}^n$, initial trust region radius $\Delta_0 > 0$ and integer $n+1 \leq p \leq (n+1)(n+2)/2$.

- 1: Build an initial interpolation set $Y_0 \subset B(\mathbf{x}_0, \Delta_0)$ of size p , with $\mathbf{x}_0 \in Y_0$.
- 2: **for** $k = 0, 1, 2, \dots$ **do**
- 3: Given \mathbf{x}_k and Y_k , construct the model $m_k(\mathbf{s})$ (2.1) satisfying (2.2).
- 4: Approximately solve the trust region subproblem (2.4) to get a step \mathbf{s}_k .
- 5: **if** $\|\mathbf{s}_k\| < \Delta_k/2$ **then**
- 6: Safety Phase: Set $\mathbf{x}_{k+1} = \mathbf{x}_k$, $\Delta_{k+1} < \Delta_k$, and form Y_{k+1} by improving the geometry of Y_k .
- 7: **goto** line 2.
- 8: **end if**
- 9: Evaluate $f(\mathbf{x}_k + \mathbf{s}_k)$ and calculate ratio r_k (2.5).
- 10: **if** $r_k \geq 0.1$ **then**
- 11: Successful Phase: Accept step $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{s}_k$ and set $\Delta_{k+1} > \Delta_k$.
- 12: **else if** *geometry of Y_k is not good* **then**
- 13: Model Improvement Phase: Reject step $\mathbf{x}_{k+1} = \mathbf{x}_k$, set $\Delta_{k+1} < \Delta_k$ and improve the geometry of Y_{k+1} .
- 14: **else**
- 15: Unsuccessful Phase: Reject step $\mathbf{x}_{k+1} = \mathbf{x}_k$ and set $\Delta_{k+1} < \Delta_k$.
- 16: **end if**
- 17: Form $Y_{k+1} = Y_k \cup \{\mathbf{x}_k + \mathbf{s}_k\} \setminus \{\mathbf{y}\}$ for some $\mathbf{y} \in Y_k$.
- 18: **end for**

The last two are designed primarily for noisy objectives.

When a restart is called, we increase the trust region radius $\Delta_{k+1} = \Delta_{\text{reset}}$ and improve the geometry of Y_k in this new trust region via two possible mechanisms:

Hard restart: Rebuild Y_k in the new trust region from scratch using the same mechanism as how Y_0 was originally constructed; and,

Soft restart (moving \mathbf{x}_k): First, move \mathbf{x}_k to a geometry-improving point in the new trust region, shifting the trust region to this new point. Then, move the $N - 1 < p$ points in Y_k which were closest to the old value of \mathbf{x}_k to geometry-improving points in the new trust region centred at the (shifted) \mathbf{x}_k . Py-BOBYQA continues from whichever of these N new points has the least objective value.


The final solution returned by the solver takes the optimal value seen so far, including the saved endpoints from previous restarts. The soft restart approach with $N = \min(3, p)$ is the default approach in Py-BOBYQA. The value of Δ_{reset} is Δ_0 by default.

Adaptive restarts The new restart set up in Py-BOBYQA, to help with (faster) escape from local minima proceeds as follows. Instead of setting Δ_{reset} to Δ_0 every time a restart is triggered (which we refer to as a fixed restart), Py-BOBYQA allows Δ_{reset} to be increased by a constant factor (default is 1.1) if the previous restart failed to produce a reduction in the objective. Thus Δ_{reset} is still set to Δ_0 for the first restart, but will be progressively increased on restarts that do not make progress.

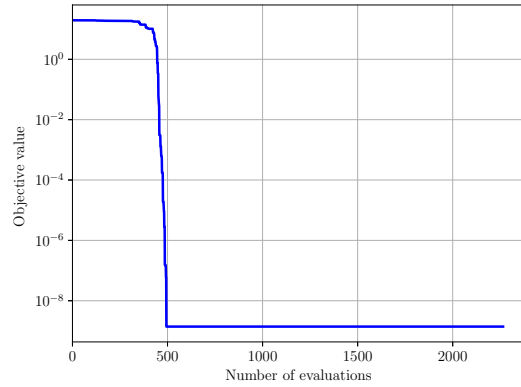
An illustrative example We use the well-known Ackley function, that has a global minimum surrounded by a high number of local minima, and illustrate the impact of using fixed and adaptive restarts in Py-BOBYQA on finding the global minimum. In Figure 1, Py-BOBYQA with $2n + 1$ initialization points is applied to the two-dimensional Ackley function with bound constraints $\mathbf{x} \in [-26, 26]^2$ and no noise, from starting point $\mathbf{x}_0 = [20, 20]^\top$ and with an evaluation budget of 10^3 gradients³. Both usual (fixed) restarts and adaptive restarts are tested. Figure 1 shows that both fixed and adaptive restarts help Py-BOBYQA reach the global minimum to high accuracy (while the no restarts Py-BOBYQA does not). Furthermore, by comparison to

³That is, a budget of $10^3(n + 1)$ evaluations for a problem of dimension n .


fixed restarts, adaptive ones both achieve a larger reduction in the cost function within fewer evaluations and require less restarts to achieve this improvement.



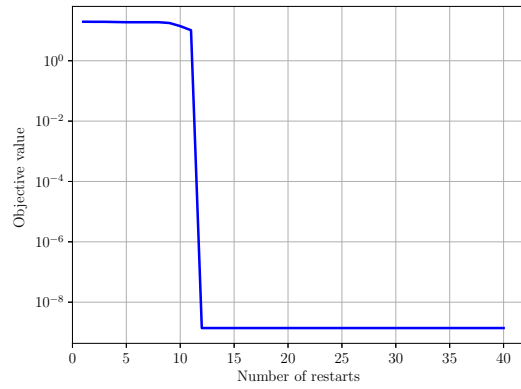
(a) Fixed restarts, evaluation budget versus objective decrease.




(b) Adaptive restarts, evaluation budget versus objective decrease.




(c) Fixed restarts, number of restarts versus objective decrease.



(d) Adaptive restarts, number of restarts versus objective decrease.



(e) Fixed restarts, path of iterates after each restart.



(f) Adaptive restarts, path of iterates after each restart.

Figure 1. Demonstration of Py-BOBYQA with multiple restarts on the Ackley function.

3 Global solvers: a brief description

Here we briefly describe the global solvers we test against local DFO. We chose a number of solvers that use global models, also constructed from just black-box, possibly noisy and expensive,

function values, compared to local models as in Py-BOBYQA, and then a small number of other state-of-the-art deterministic and stochastic global solvers. Our aim here is not to provide a comprehensive benchmarking exercise for global optimization solvers, but to make sure we choose state-of-the-art global solvers with a similar remit, to compare with Py-BOBYQA. We refer the interested reader to the extensive comparisons and clear descriptions of state-of-the-art derivative-free local and global solvers given in [37], and earlier in [1], which we also used as a guide for some of our solver choices.

Bayesian optimization solvers Ever since [28, 23], Bayesian methods have been designed for the global optimization of black-box functions, which are typically expensive to compute. This remit implies that the methods spend computational time and resources to make informed decisions about where to sample the function in the course of their runs. As such, Bayesian algorithms attempt to strike a balance between *exploring* unsampled regions, and *exploiting* regions with a promising objective value; see [5, 39, 12] and the references therein for detailed surveys of this area and connections to global optimization methods and applications.

As the name indicates, these methods rely on Bayes’ theorem and related Bayesian processes in probability, namely, they assume a (global) model (a prior distribution) of the objective f is available together with some point samples from this distribution, say \mathcal{D}_k at iteration k of the algorithm, which are iteratively used to determine an a posteriori distribution for f and evolve the current model for f . Typical choices of prior distributions for the objective include Gaussian processes (used in the GPyOpt solver [41]), random forests (PySMAC [17]), regression trees, and tree parzen estimators (HyperOpt [3]); see [35] for more details and descriptions. The prior/current model of f implies (often closed-form) expressions or choices for the sampled mean $\mu(\cdot|\mathcal{D}_k)$ and covariance/noise kernel $\sigma(\cdot|\mathcal{D}_k)$. The noise kernel allows a Gaussian process to capture a wide range of functional/distributional constructions; various options have been proposed, such as the well-known squared exponential/Gaussian kernel, Matérn ν -kernel, radial basis functions (used in PySOT [11]), and more [35].

Using and combining the expectation and covariance (noise kernel) of the distribution, an acquisition function $a(\cdot)$ is formed, which is then employed in order to find the next sampling point, $\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} a(\mathbf{x})$. Instead of minimizing directly the model of the objective, as it is done in local DFO for example, Bayesian optimization methods minimize some measure of ‘error’ or ‘discrepancy’ between the objective and its model. The latter optimization problem is simpler, and it is inexpensive with respect to evaluations of the objective. The choice and optimization of the acquisition function determines whether the next point \mathbf{x}_{k+1} /model leads to exploring areas with higher uncertainty, or exploiting regions with more promising expected values. Various choices of acquisition functions are in use, such as the *expected improvement*, *probability of improvement*, *upper confidence bounds* and more. The acquisition function is updated with the new point that was calculated and its associated function value, iteratively, until convergence or the given budget are reached. Then, in addition to optimization solution information, the acquisition function can be extracted, which contains a distribution and confidence levels for the solution.

Bayesian optimization is frequently used in machine learning applications, such as for hyperparameter tuning of optimization algorithms for training classifiers [18, 40, 26]. Ghanbari & Scheinberg [13] compared derivative-free local solvers against Bayesian global solvers on such hyperparameter tuning problems, with encouraging results.

A surrogate optimization algorithm Surrogate or response surface optimization methods are similar to Bayesian algorithms [14, 12, 21], in that they use a global model/surrogate (such as a radial basis function) to approximate the black-box objective function and then iteratively refine this model. This model is deterministic and constructed by interpolation of function values; a ‘bumpiness’ function measures the discrepancy between the model and true function, and its minimization is used to find the next interpolation point (similarly to the acquisition function in Bayesian methods). Gaussian priors can be viewed as particular radial basis functions, and so the two classes of methods are closely connected [12]. PySOT [11] is a state-of-the-art surrogate optimization solver that uses radial basis functions for the response surface; we have included it

here due to its similarities to Bayesian optimization, in order to assess the potential benefits of different global models.

A branch and fit method SNOBFIT [19] constructs local quadratic models from interpolation of function values, and uses a branching strategy to enhance their global properties. It connects well with Py-BOBYQA in that it is a local method, with improvements that make it suitable for global optimization, hence our choice to include it here.

An evolutionary algorithm Covariance Matrix Adaptation evolutionary strategy (CMA-ES) [15, 16] is a genetic algorithm, that attempts to find the basin of a global minimizer by iteratively moving and reshaping a Gaussian distribution using samples of objective values (populations). The point with the best objective value is labelled as the ‘best point’ and the mean of the distribution is updated, which shifts the mean of the distribution in the direction of the best point. Updating the covariance of the distribution determines its shape and spread, with the algorithm deciding how much old information to keep and how much to update from current best point obtained. In CMA-ES, the updating of the kernel/covariance is a weighted formula allowing three types of contributions (past, rank-one updates along a direction of progress or higher rank updates); the weights in the covariance updating formula are also adjusted in the algorithm. The sophisticated updates are key to the success of this evolutionary strategy and can be viewed as providing some curvature information for the objective.

The DIRECT algorithm [22] The well-known DIvide a hyperRECTangle (DIRECT) method is an extension of Lipschitz global optimization strategies that does not require the Lipschitz constant to be known/provided, but uses trial values of this constant chosen heuristically [14]. In particular, the domain is subdivided into hyper-rectangles and the function is evaluated at the midpoint of each box, and this evaluation is then used to construct Lipschitz-like lower bounds. The splitting rules for subdomains are carefully chosen to focus on those that yield the largest potential decrease in the objective, so as to encourage the algorithm towards global searches.

4 Testing methodology

Solvers tested We compared Py-BOBYQA with a selection of global optimization routines, all implemented or wrapped in Python⁴. Here, we list each solver tested and any non-standard parameters set—in some cases, these choices depended on whether we were testing an objective with or without noise:

CMA-ES v2.5.7 [15, 30]: Used an initial standard deviation $0.2 \min_i(b_i - a_i)$, where the feasible region is $\mathbf{a} \leq \mathbf{x} \leq \mathbf{b}$;

DIRECT v1.0.1 [22, 20]: Used all defaults, except for increasing the maximum number of objective evaluations to be sufficiently large for our testing;

GPyOpt v1.2.1 [41]: Used an initial design of $2n$ random points, and set the `exact_feval` flag depending on whether noisy is objective or not;

HyperOpt v0.1 [3, 4]: Used a uniform search space over the feasible region, Tree of Parzen Estimators algorithm, with and without providing noise variance for the objective evaluation;

PySMAC v0.9.1 [17, 2]: We set the `deterministic` option according to whether the objective had noise or not;

PySOT v0.1.36 [11]: Used a radial basis function surrogate with an initial design of $2n + 1$ Latin Hypercube points and a DYCORS [36] strategy with $100n$ points for selecting new points;

SNOBFIT v1.0.0 [19, 43]: Used all default parameters.

⁴We used Python 2.7 with NumPy 1.12.1 and SciPy 1.0.1.

For Py-BOBYQA, we used version 1.1⁵, with options: scale feasible region to $[0, 1]^n$; initial trust region radius 0.1 and final trust region radius 10^{-8} (in scaled variables); on unsuccessful restarts, increase trust region radius Δ_{reset} by a factor⁶ of 1.1; terminate after 10 consecutive unsuccessful restarts or 20 total unsuccessful restarts.

All solvers were allowed a budget of at most $10^4(n + 1)$ objective evaluations for an n -dimensional problem, and a runtime of at most 12 hours per problem instance. In addition, all solvers except DIRECT require a starting point; for each solver, problem, and instance, this was chosen independently and uniformly from the entire feasible region. Given this, we ran 10 instances of each solver on each problem (where, if applicable, each instance had a different starting point, and realizations of the stochastic noise in the objective).

Performance measurement Following the approach from Moré & Wild [29], we measure solver performance by counting the number evaluations required to first achieve a sufficient objective reduction; that is, for a solver \mathcal{S} , problem p (with true objective f and noisy realized objective \tilde{f}), and accuracy level $\tau \in (0, 1)$, we define the number of evaluations required to solve the problem as:

$$N_p(\mathcal{S}; \tau) := \# \text{ objective evaluations taken to find a point } \mathbf{x} \text{ satisfying} \\ \tilde{f}(\mathbf{x}) \leq \mathbb{E}[\tilde{f}(\mathbf{x}^*) + \tau(\tilde{f}(\mathbf{x}_0) - \tilde{f}(\mathbf{x}^*))], \quad (4.1)$$

where \mathbf{x}^* is the global minimizer of f (also the minimizer of $\mathbb{E}[\tilde{f}]$). If (4.1) was not achieved at all by the solver, we set $N_p(\mathcal{S}; \tau) = \infty$. We also note that $\tilde{f}(\mathbf{x}_0)$ is different for each instance of each problem (and each solver).

To compare solvers, we will use data profiles [29]⁷. Data profiles measure the proportion of problems solved within a fixed number of evaluations (normalized by problem dimension); that is, for each solver \mathcal{S} on a set of test problems \mathcal{P} , we plot the curve

$$d_{\mathcal{S}}(\alpha) := \frac{1}{|\mathcal{P}|} \cdot |\{p \in \mathcal{P} : N_p(\mathcal{S}; \tau_p) \leq \alpha(n_p + 1)\}|, \quad \text{for } \alpha \geq 0, \quad (4.2)$$

where n_p is the dimension of problem p . In our testing, we ran each solver on 10 instances of each problem (i.e. independent realizations of the starting point and/or random noise in objective evaluations). For data profiles, we take each problem instance as a separate problem to be ‘solved’; i.e. we plot the proportion of the 500 problem instances solved within a given computational budget.

For noisy problems, some of the objective reduction in (4.1) may be due to a ‘lucky’ realization of the noise, rather than genuine progress towards the objective. To ensure that our results only capture genuine progress in the optimization, we follow [6] and floor τ on a per-problem basis; specifically, for each solve instance of each problem we define:

$$\tau_{\text{crit}}(p) := 10^{\lceil \log_{10} \hat{\tau}(p) \rceil}, \quad \text{where} \quad \hat{\tau}(p) := \frac{\sigma(\mathbf{x}^*)}{\mathbb{E}[\tilde{f}(\mathbf{x}_0) - \tilde{f}(\mathbf{x}^*)]}, \quad (4.3)$$

where $\sigma(\mathbf{x}^*)$ is the standard deviation of $\tilde{f}(\mathbf{x}^*)$. We then replace τ in (4.1) with

$$\tau_p := \min(\tau_{\text{max}}, \max(\tau_{\text{crit}}(p), \tau)), \quad (4.4)$$

where $\tau_{\text{max}} = 10^{-1}$ and τ is the actual accuracy level desired (e.g. $\tau = 10^{-5}$).

⁵Available from <https://github.com/numericalalgorithmsgroup/pybobyqa>

⁶The constant factor 1.1 was not decided based on testing or tuning; it was simply considered to be a sensible choice.

⁷Another possibility is to use performance profiles [29], where the ‘difficulty’ of the problem is also taken into account. We have generated these as well, and compared them with data profiles. We found that they did not bring new/different information in our tests compared to data profiles and so we have not included them here for space reasons.

Lastly, we note that DIRECT does not allow the specification of a starting point, and instead always selects \mathbf{x}_0 to be the center of the feasible region. For 13 problems⁸, this point is the global minimizer, and so DIRECT appears to solve these problems to global optimality after exactly one objective evaluation. To provide a more reasonable measure of DIRECT’s performance, in these instances, we replace \mathbf{x}_0 by \mathbf{x}_1 in (4.1) and look for the second time this reduction is achieved (i.e. excluding \mathbf{x}_0 from the count).

4.1 Test problems

GO test set For our comparison, we used the collection of 50 problems from Ali, Khompatraporn, and Zabinsky [1]. To make the testing more difficult, we increased the dimension of 7 problems to $n \in [25, 50]$, where they had $n = 10$ in [1]. The collection of test problems is given in Table 1 in Appendix A.

In addition, we optionally modified the problems by the addition stochastic noise; that is, each solver has access to $\tilde{f}(\mathbf{x}) \approx f(\mathbf{x})$ for each \mathbf{x} . We tested the following noise models:

- Smooth (no noise): $\tilde{f}(\mathbf{x}) = f(\mathbf{x})$
- Multiplicative Gaussian noise: $\tilde{f}(\mathbf{x}) = f(\mathbf{x})(1 + \sigma\epsilon)$, where $\epsilon \sim N(0, 1)$ i.i.d. for each \mathbf{x} and each solver; and
- Additive Gaussian noise: $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \sigma\epsilon$, where $\epsilon \sim N(0, 1)$ i.i.d. for each \mathbf{x} and each solver.

For our testing, we used a noise level of $\sigma = 10^{-2}$.

Hyperparameter tuning (MNIST) We also considered a case study of hyperparameter tuning for machine learning, for which developers have often used Bayesian optimization solvers. Here, we followed the approach from Snoek, Larochelle, and Adams [40], and consider the problem of training a multi-class logistic classifier for MNIST⁹ using stochastic gradient descent with a fixed learning rate. The 4 hyperparameters we optimized in our testing were:

- \log_{10} of learning rate, in $[-10, 0]$;
- \log_{10} of ℓ_2 regularization parameter, in $[-10, 0]$;
- Number of epochs, in $[1, 100]$; and
- Batch size, in $[1, 2000]$.

We converted real-valued inputs for the last two parameters to integer values by rounding up to the nearest integer. The objective function we selected was the (top-1) error rate on the test set. We ran the solvers as per Section 4 (i.e. selecting a random feasible \mathbf{x}_0 if needed), for a maximum of $100(n + 1) = 500$ objective evaluations or 36 hours. For each solver, we ran 30 instances of the problem, and compared performance using data and performance profiles, in the same way as above.

Selecting the ‘optimum’ value for MNIST: To construct data and performance profiles, we need estimates of $\mathbb{E}[\tilde{f}(\mathbf{x}_0)]$, $\mathbb{E}[\tilde{f}(\mathbf{x}^*)]$ and $\sigma(\mathbf{x}^*)$ in order to evaluate (4.1). However, unlike before, we do not have access to the true values for our objective, and so we need to provide sensible estimates of these values. For $\mathbb{E}[\tilde{f}(\mathbf{x}_0)]$, we used the value of $\tilde{f}(\mathbf{x}_0)$ from each instance of each solver. To determine \mathbf{x}^* , and hence $\mathbb{E}[\tilde{f}(\mathbf{x}^*)]$ and $\sigma(\mathbf{x}^*)$, we did the following:

1. Take the 20 points found by all instances of all solvers with the smallest (noisy) objective value;

⁸These are: Ackley, Bohachevsky 1, Bohachevsky 2, Camel 3, Cosine Mixture, Exponential, Griewank, Periodic, Powell Quadratic, Rastrigin, Salomon, Schaffer 1, and Schaffer 2.

⁹The multi-class classifier came from by training 10 one-versus-all classifiers, with prediction by selecting the class with highest probability. The training set had 60,000 images, the test set had 10,000 images.

2. Evaluate the (noisy) objective \tilde{f} independently 30 times at each of these points \mathbf{x} , and use this to estimate $\mathbb{E}[\tilde{f}(\mathbf{x})]$ and $\sigma(\mathbf{x})$;
3. Select as \mathbf{x}^* the point which gave the smallest value of $\mathbb{E}[\tilde{f}(\mathbf{x})]$.

This process gave the point $\mathbf{x}^* = [-5.2706, -2.2505, 100, 411.02]^\top$, with estimates $\mathbb{E}[\tilde{f}(\mathbf{x}^*)] = 0.0818067$ and $\sigma(\mathbf{x}^*) = 1.01388 \times 10^{-3}$. This equates to a learning rate of 5.36×10^{-6} , ℓ_2 regularization parameter 5.62×10^{-3} , 100 epochs and batch size 412, yielding a test set error rate of $8.1\% \pm 0.1\%$.

5 Numerical results

5.1 Numerical results on the GO test set

Selecting the best-performing Py-BOBYQA variant We compare various Py-BOBYQA variants on the GO test set, to find which variant is most efficient for different accuracy/budget regimes. We divided our testing based on the size of the interpolation set, namely, $p = 2n + 1$ and $p = (n + 1)(n + 2)/2$, which corresponds to the case of an inexact local quadratic (as if having a diagonal Hessian) and a full quadratic model m_k in (2.1), respectively¹⁰. Clearly, building and using an inexact quadratic local model has a reduced evaluation and computational cost (compared to a full quadratic model), but we need to assess whether the local efficiency translates into globally good performance, across iterations; hence the need to look at the two types of local models. For both choices of interpolation models, we compare the different types of restarts allowed in Py-BOBYQA—*hard* and *soft* (choice of interpolation set at a restart); *fixed* and *adaptive* (size of trust region on restart).

Low accuracy/budget ($\tau = 10^{-2}$ in (4.1) and budget 10^2 gradients) Our tests show that in this accuracy/budget regime, there is no substantial difference between Py-BOBYQA variants, for the different interpolation models and restarts strategies. We note a (small but clear) reduction in performance when no restarts are performed, and a (slightly) superior performance for soft restarts with fixed and adaptive choices of trust region radius on restarts, for both values of the interpolation set size p . We also note that the evaluation cost of building full quadratic local models can be seen in the early iterations of the algorithms, when the variants that use reduced quadratics can already make progress, while the full size models are still being built; this can be seen in two cases in Figure 9. We delegate the data profiles for this case to Appendix B, where we illustrate also the different noise models and compare the two interpolation cases.

High accuracy/high budget ($\tau = 10^{-5}$ in (4.1) and budget of 10^4 gradients) For both interpolation models, and different noise models (including no noise), Figures 2–4 show the continuously improving performance of some Py-BOBYQA variants with restarts, while the no-restarts variant for example, can no longer improve its performance beyond the achievements in the low accuracy regime, despite the substantially larger budget. In particular, the no-restarts Py-BOBYQA is the worst performing variant, failing to solve the largest amount of problems within the given (large) budget. In terms of restarts, the soft restart variant with adaptive restart radius (adaptive Δ_0) is the most efficient and economical across all models and restart types, followed by standard soft restarts (with fixed radius size) and matched sometimes by hard restarts with adaptive radius size (such as for $p = 2n + 1$ and noisy problems; and $p = (n + 1)(n + 2)/2$ and no noise). To contrast the two interpolation model types, we see that Py-BOBYQA with full quadratic models ($p = (n + 1)(n + 2)/2$), soft restarts and adaptive radius on restarts gives the best performance overall in our GO tests, and is particularly effective on noisy problems. We note that the superior performance of this Py-BOBYQA variant is matched by Py-BOBYQA with inexact quadratic models ($p = 2n + 1$) and similar soft and adaptive restarts on smooth problems, which is important as the computational cost of Py-BOBYQA ($p = 2n + 1$) is less than that of Py-BOBYQA with full quadratic models, as the linear interpolation system that is solved to find

¹⁰The case $p = n + 1$ when m_k is linear (and hence contains no curvature information) does not yield competitive variants for general objectives [9, 7].

the local model is much smaller in dimension. We also note that in the large budget regime, we no longer ‘pay the price’ of the more expensive model set up.

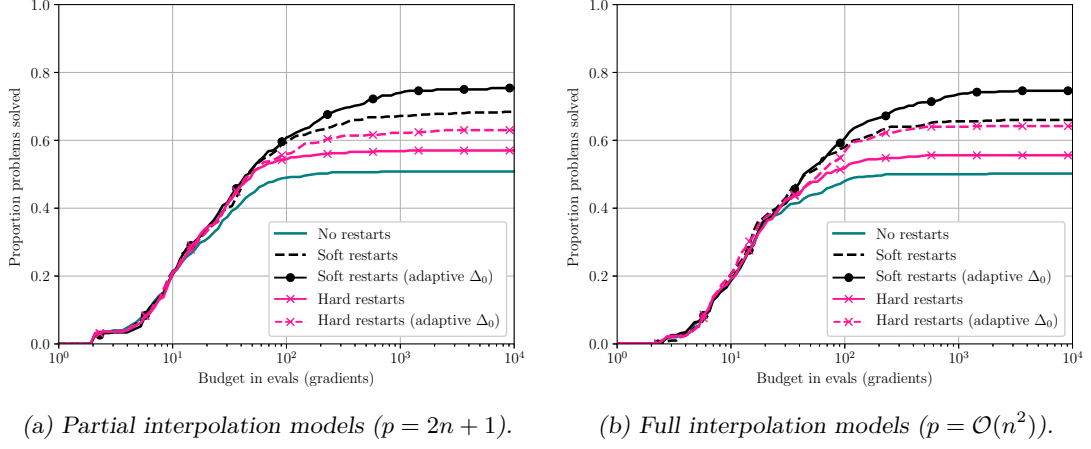


Figure 2. Comparison of Py-BOBYQA variants on the smooth GO test set.




Figure 3. Comparison of Py-BOBYQA variants on the GO test set with multiplicative noise.

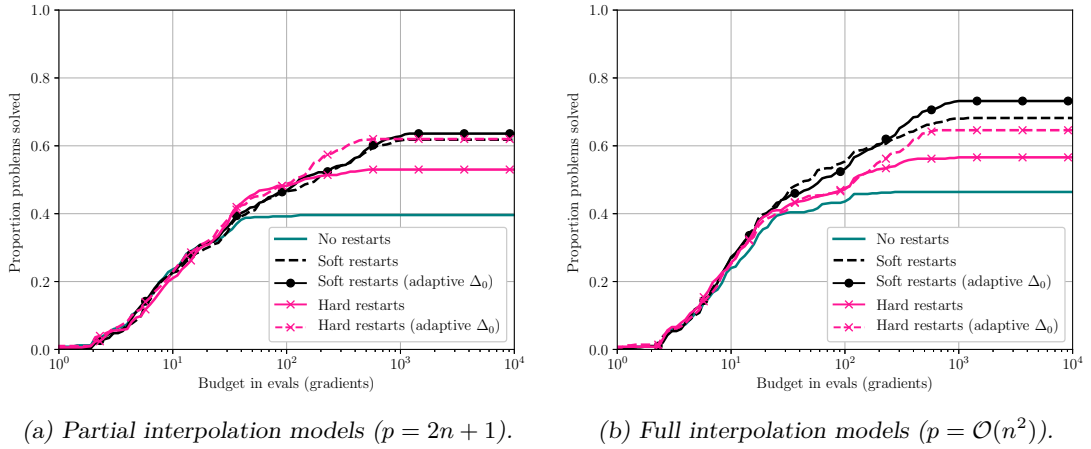


Figure 4. Comparison of Py-BOBYQA variants on GO test set with additive noise.

Py-BOBYQA variance test As Py-BOBYQA is a local solver, and as we are using (ten) randomly chosen starting points per problem instance, we investigate here the variation in performance over ten different runs of Py-BOBYQA over the test set with different starting points each time. This illustrates the variance in our previous results, and how much the performance is influenced by a particular starting point, a key question for local solvers. Given the good performance of Py-BOBYQA with full quadratic models ($p = \mathcal{O}(n^2)$) and soft restarts across all test problem variants, we are focusing on this variant, with fixed and adaptive restart trust-region radius on restarts. We are only looking at smooth problems (with no noise), to make sure the only randomness present is due to choice of starting point. In Figure 5, we show separate data profiles for the ten runs of Py-BOBYQA (i.e. each profile shows one instance of each problem with a random starting point), and the average data profile.

The variation in starting point usually leads to between 5%–10% variation in performance in the second half/end of given budget, except in the low accuracy regime and adaptive trust region radius where we see a variation in performance of about 20% at the end of the budget. For high accuracy/large budget, the variation in performance is such that if we compare with Figure 2, we see that all/most adaptive radius runs are better in performance than the fixed radius variant. For low accuracy, the situation is less clear and performances of different variants are less distinctive, which is no surprise given the average plots in Figure 9.

We also note that these conclusions are similar for Py-BOBYQA with $p = 2n + 1$ interpolation points; see Figure 10 in Appendix B.

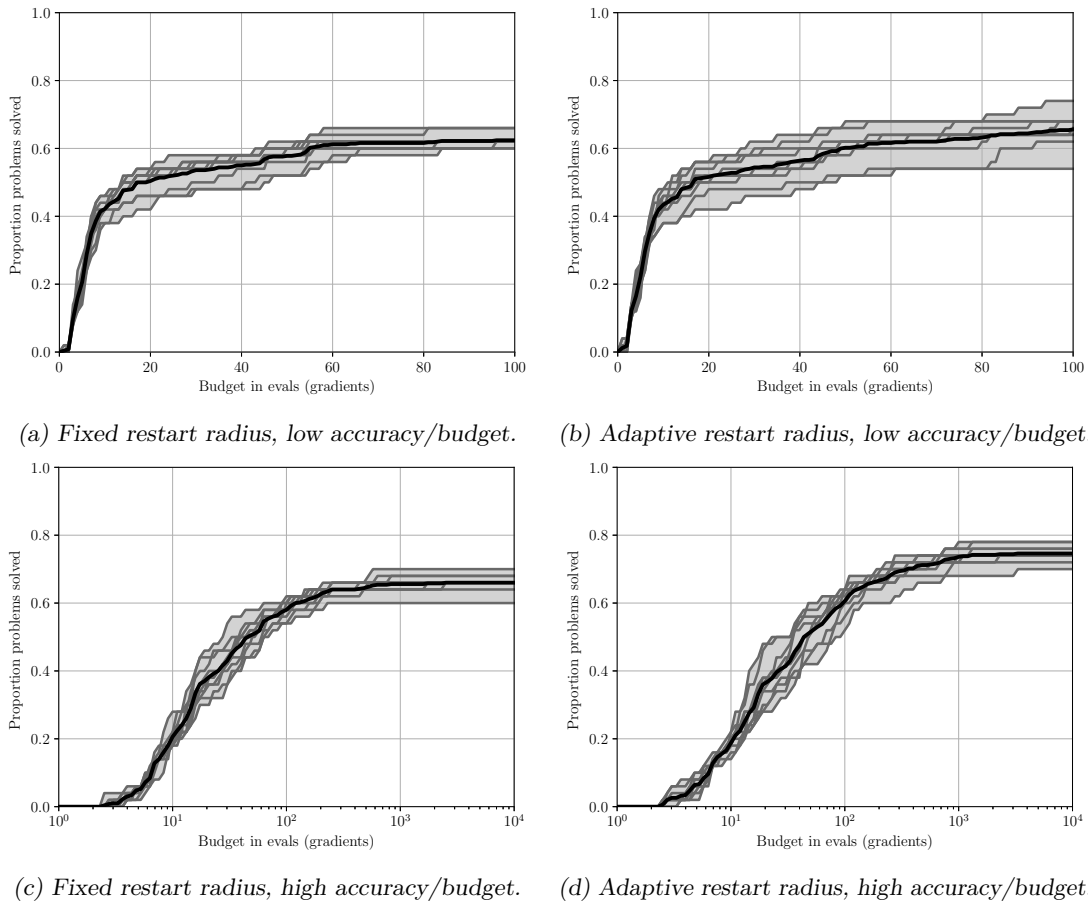


Figure 5. Demonstration of impact of random starting point for Py-BOBYQA ($p = \mathcal{O}(n^2)$ with soft restarts). Low accuracy and budget are $\tau = 10^{-2}$ and 10^2 gradients respectively; high accuracy and budget are $\tau = 10^{-5}$ and 10^4 gradients respectively. The dark lines are the average of all runs.

Selecting the best-performing Bayesian and surrogate solver We first compare the Bayesian solvers (GPyOpt, HyperOpt and PySMAC) and surrogate algorithm PySOT for smooth and noisy GO test set.

Figure 6 shows that, for smooth problems, PySOT is consistently the best performing solver compared to all Bayesian solvers that we tested, in both low and high accuracy regimes; the next best solver is GPyOpt for low accuracy regime, but when large budgets are available, HyperOpt and PySMAC improve their performance (compared to GPyOpt). The same results apply for noisy problems; these results are shown in Figure 11 (Appendix B). Thus in our next set of comparisons between all solvers, we select only PySOT to ‘represent’ the best behaviour of Bayesian and surrogate solvers.

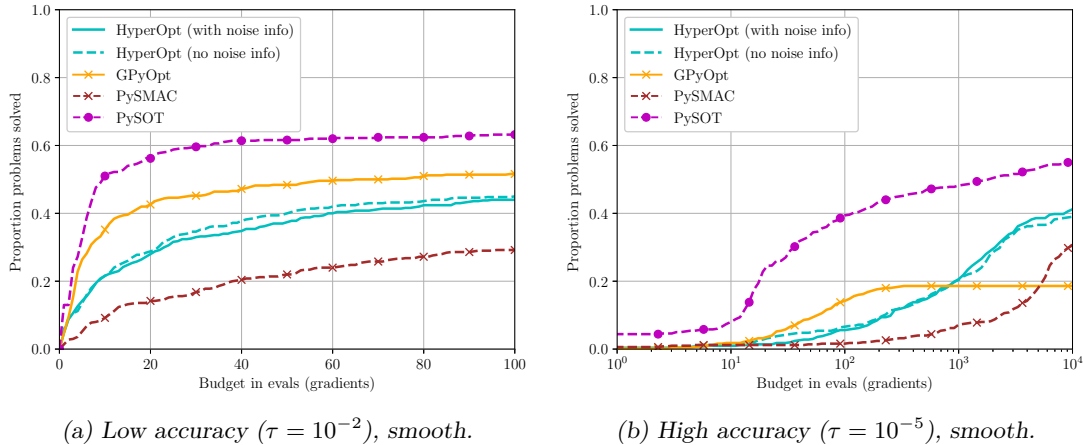


Figure 6. Comparison of Bayesian and surrogate solvers on the GO test set (smooth case; see Figure 11 in Appendix B for noisy results).


Comparisons of all solvers In Figure 7, we compare Py-BOBYQA against the global solvers; ‘‘Py-BOBYQA best’’ represents the variant with full quadratic interpolation model ($p = \mathcal{O}(n^2)$), soft restarts and adaptive trust region radius on restarts, that we found earlier to be performing best on the GO test set amongst all Py-BOBYQA variants, in all regimes of accuracy and noise; see Figures 2–4, Figure 9, and our corresponding comments above. Also, ‘‘PySOT best’’ represents the best performing solver on the GO test set amongst the chosen Bayesian and surrogate solvers.

Low accuracy/budget ($\tau = 10^{-2}$ in (4.1) and budget 10^2 gradients) The first column in Figure 7 shows that on average, DIRECT performs best in this regime, and the remaining solvers are essentially similar at the end of the budget. In the earlier phases of budget, Py-SOT and then the best Py-BOBYQA variant have best performance, being able to solve more problems on small budgets than the remaining solvers, while CMA-ES and SNOBFIT seem to need longer budgets to reach a similar amount of solved problems at the end of the low accuracy budget.


High accuracy/budget ($\tau = 10^{-5}$ in (4.1) and budget 10^4 gradients) For smooth problems, the best Py-BOBYQA variant is the best performing solver followed by DIRECT, while for multiplicative noise, these two solvers have comparably good performance. For additive noise, DIRECT has superior performance, followed by similarly good performance from PySOT and Py-BOBYQA.

5.2 Numerical results on MNIST


We recall the description of the hyper-parameter tuning problem for the MNIST test set in Section 4, where we described the four hyper-parameters, the objective value and the low accuracy level that is needed here. Again, we compare first the Bayesian and surrogate optimization solvers in Figure 8a; we find similar performance to the low accuracy regime for the GO test set (Figure 6), namely, the solvers are comparable with PySOT being slightly better. When comparing the remaining solvers with the best PyBOBYQA and PySOT in Figure 8b, we find that DIRECT




(a) Low accuracy ($\tau = 10^{-2}$), smooth.




(b) High accuracy ($\tau = 10^{-5}$), smooth.




(c) Low accuracy ($\tau = 10^{-2}$), multiplicative noise.



(d) High accuracy ($\tau = 10^{-5}$), multiplicative noise.



(e) Low accuracy ($\tau = 10^{-2}$), additive noise.




(f) High accuracy ($\tau = 10^{-5}$), additive noise.

Figure 7. Comparison of all and best-performing solvers on the GO test set.

and PySOT are best performing, and the best performing Py-BOBYQA variant is not much worse.

6 Conclusions and future directions

We investigated numerically the behaviour of a local derivative-free model-based trust region method Py-BOBYQA, and a new variant, on global optimization problems. We found encouraging performance of the improved variant when compared to Bayesian and surrogate algorithms, with



(a) Bayesian and surrogate solvers.

(b) All/best performing solvers.

Figure 8. Comparison of solvers on the MNIST problem, low accuracy ($\tau = 10^{-2}$).

DIRECT, SNOBFIT and CMA-ES solvers, for both smooth and noisy problems solved to low or high accuracy requirements. Our experiments illustrate that local derivative free methods may be effective for global optimization, in the regime typical for global solvers, namely, when the objective is black-box, possibly noisy and computationally expensive to evaluate. This finding is useful as these local methods have greater flexibility (such as allowing the use of a good initial guess of the solution), guarantee improvement in current best guess of the solution, and have better scalability. However, we delegate to future directions the crucial task of estimating the performance of the local DFO solver on large-scale problems, both in terms of computational time and evaluations. Theoretical developments that would show that the improved Py-BOBYQA can escape local minima would also be desirable.

6.1 Acknowledgements

We acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility¹¹ in carrying out this work.

References

- [1] M. M. ALI, C. KHOMPATRAPORN, AND Z. B. ZABINSKY, *A numerical evaluation of several stochastic algorithms on selected continuous global optimization test problems*, J. Glob. Optim., 31 (2005), pp. 635–672.
- [2] AUTOML, *pySMAC*. <https://github.com/automl/pysmac>, 2018.
- [3] J. BERGSTRA, D. L. K. YAMINS, AND D. D. COX, *Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures*, in 30th Int. Conf. Mach. Learn., Atlanta, 2013, pp. 115–123.
- [4] ———, *Hyperopt: Distributed Asynchronous Hyper-parameter Optimization*. <https://github.com/hyperopt/hyperopt>, 2018.
- [5] E. BROCHU, V. M. CORA, AND N. DE FREITAS, *A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning*, arXiv preprint arXiv:1012.2599, (2010).
- [6] C. CARTIS, J. FIALA, B. MARTEAU, AND L. ROBERTS, *Improving the flexibility and robustness of model-based derivative-free optimization solvers*, tech. rep., University of

¹¹<http://dx.doi.org/10.5281/zenodo.22558>

- Oxford, Mathematical Institute, 2018. Available on Optimization Online, http://www.optimization-online.org/DB_HTML/2018/03/6549.html.
- [7] C. CARTIS AND L. ROBERTS, *A derivative-free Gauss-Newton method*, tech. rep., University of Oxford, Mathematical Institute, 2017. Available on Optimization Online, http://www.optimization-online.org/DB_HTML/2017/10/6304.html.
- [8] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *Trust-Region Methods*, MPS-SIAM Series on Optimization, MPS/SIAM, Philadelphia, 2000.
- [9] A. R. CONN, K. SCHEINBERG, AND L. N. VICENTE, *Introduction to Derivative-Free Optimization*, vol. 8 of MPS-SIAM Series on Optimization, MPS/SIAM, Philadelphia, 2009.
- [10] A. L. CUSTÓDIO, K. SCHEINBERG, AND L. N. VICENTE, *Methodologies and Software for Derivative-free Optimization*, in Adv. Trends Optim. with Eng. Appl., T. Terlaky, M. F. Anjos, and S. Ahmed, eds., MOS-SIAM Book Series on Optimization, SIAM, Philadelphia, 2017.
- [11] D. ERIKSSON, D. BINDEL, AND C. SHOEMAKER, *Surrogate optimization toolbox (pySOT)*. <https://github.com/dme65/pySOT>, 2015.
- [12] J. FOWKES, *Bayesian Numerical Analysis: Global Optimization and Other Applications*, PhD thesis, Oxford University, 2011.
- [13] H. GHANBARI AND K. SCHEINBERG, *Black-box optimization in machine learning with trust region based derivative free algorithm*, arXiv preprint arXiv:1703.06925, (2017).
- [14] H.-M. GUTMANN, *Radial Basis Function Methods for Global Optimization*, PhD thesis, Cambridge University, 2001.
- [15] N. HANSEN AND A. OSTERMEIER, *Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation*, in Proc. IEEE Int. Conf. Evol. Comput., Nagoya, 1996, pp. 312–317.
- [16] ———, *Completely derandomized self-adaptation in evolution strategies*, Evolutionary Computation, 9 (2001), pp. 159–195.
- [17] F. HUTTER, H. H. HOOS, AND K. LEYTON-BROWN, *Sequential Model-Based Optimization for General Algorithm Configuration*, in Int. Conf. Learn. Intell. Optim., Rome, 2011, pp. 507–223.
- [18] F. HUTTER, H. H. HOOS, AND T. STÜTZLE, *Automatic algorithm configuration based on local search*, in AAAI, vol. 7, 2007, pp. 1152–1157.
- [19] W. HUYER AND A. NEUMAIER, *SNOBFIT – Stable Noisy Optimization by Branch and Fit*, ACM Trans. Math. Softw., 35 (2008), pp. 1–25.
- [20] J. M. GABLONSKY, *pydirect*. <https://code.google.com/archive/p/pydirect>, 2012.
- [21] D. R. JONES, *A taxonomy of global optimization methods based on response surfaces*, Journal of global optimization, 21 (2001), pp. 345–383.
- [22] D. R. JONES, C. D. PERTTUNEN, AND B. E. STUCKMAN, *Lipschitzian optimization without the Lipschitz constant*, J. Optim. Theory Appl., 79 (1993), pp. 157–181.
- [23] D. R. JONES, M. SCHONLAU, AND W. J. WELCH, *Efficient global optimization of expensive black-box functions*, Journal of Global optimization, 13 (1998), pp. 455–492.
- [24] C. KELLEY, *Implicit Filtering*, SIAM, 2011.
- [25] T. KOLDA, R. LEWIS, AND V. TORCZON, *Optimization by direct search: new perspectives on some classical and modern methods*, SIAM Review, 45 (2003), pp. 385–482.

- [26] L. LI, K. JAMIESON, G. DESALVO, A. ROSTAMIZADEH, AND A. TALWALKAR, *Hyperband: A novel bandit-based approach to hyperparameter optimization*, arXiv preprint arXiv:1603.06560, (2016).
- [27] M. LOCATELLI AND F. SCHOEN, *Global Optimization: Theory, Algorithms and Application*, MOS-SIAM Series on Optimization, 2013.
- [28] J. MOCKUS, V. TIESIS, AND A. ZILINSKAS, *Bayesian methods for seeking the extremum*, *Toward Global Optimization*, 2 (1978), pp. 117–129.
- [29] J. J. MORÉ AND S. M. WILD, *Benchmarking derivative-free optimization algorithms*, *SIAM Journal on Optimization*, 20 (2009), pp. 172–191.
- [30] N. HANSEN, *pycma*. <https://github.com/CMA-ES/pycma>, 2018.
- [31] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering, Springer, New York, 2nd ed., 2006.
- [32] M. J. D. POWELL, *Direct search algorithms for optimization calculations*, *Acta Numerica*, 7 (1998), pp. 287–336.
- [33] ———, *Developments of NEWUOA for minimization without derivatives*, Tech. Rep. DAMTP 2007/NA05, University of Cambridge, 2007.
- [34] ———, *The BOBYQA algorithm for bound constrained optimization without derivatives*, Tech. Rep. DAMTP 2009/NA06, University of Cambridge, 2009.
- [35] C. E. RASMUSSEN AND C. K. WILLIAMS, *Gaussian processes for machine learning. 2006*, The MIT Press, Cambridge, MA, USA, 38 (2006), pp. 715–719.
- [36] R. G. REGIS AND C. A. SHOEMAKER, *Combining radial basis function surrogates and dynamic coordinate search in high-dimensional expensive black-box optimization*, *Eng. Optim.*, 45 (2013), pp. 529–555.
- [37] L. M. RIOS AND N. V. SAHINIDIS, *Derivative-free optimization: a review of algorithms and comparison of software implementations*, *J. Glob. Optim.*, 56 (2013), pp. 1247 – 1293.
- [38] L. ROBERTS, *Derivative-free optimisation for data fitting*, Tech. Rep. InFoMM CDT Report, University of Oxford, 2016. http://people.maths.ox.ac.uk/roberts1/docs/DF0_MiniprojectReport_updateNov18.pdf.
- [39] B. SHAHRIARI, K. SWERSKY, Z. WANG, R. P. ADAMS, AND N. DE FREITAS, *Taking the human out of the loop: A review of bayesian optimization*, *Proceedings of the IEEE*, 104 (2016), pp. 148–175.
- [40] J. SNOEK, H. LAROCHELLE, AND R. P. ADAMS, *Practical Bayesian Optimization of Machine Learning Algorithms.*, in *Adv. Neural Inf. Process. Syst.*, 2012, pp. 2951–2959.
- [41] THE GPYOPT AUTHORS, *GPyOpt: A Bayesian optimization framework in Python*. <http://github.com/SheffieldML/GPyOpt>, 2016.
- [42] M. H. WRIGHT, *Direct search methods: once scorned, now respectable*, in *Numerical Analysis 1995 (Proceedings of the 1995 Dundee Biennial Conference in Numerical Analysis)*, D. Griffiths and G. Watson, eds., 1996, pp. 191–208.
- [43] Z. FU, *Package snobfit*. <http://reflectometry.org/danse/docs/snobfit>, 2009.


A Test Problems

#	Problem	Dimension	Bounds	Global minimum
1	Ackley	10	$-30 \leq x_i \leq 30$	0
2	Aluffi-Pentini	2	$-10 \leq x_i \leq 10$	-0.3523861
3	Becker and Lago	2	$-10 \leq x_i \leq 10$	0
4	Bohachevsky 1	2	$-50 \leq x_i \leq 50$	0
5	Bohachevsky 2	2	$-50 \leq x_i \leq 50$	0
6	Branin	2	$[-5, 0] \leq \mathbf{x} \leq [10, 15]$	3.926991
7	Camel 3	2	$-5 \leq x_i \leq 5$	0
8	Camel 6	2	$-5 \leq x_i \leq 5$	-1.031628
9	Cosine Mixture	4	$-1 \leq x_i \leq 1$	-0.4
10	Dekkers and Aarts	2	$-20 \leq x_i \leq 20$	-2.477289×10^4
11	Easom	2	$-10 \leq x_i \leq 10$	-1
12	Epistatic Michalewicz	5	$0 \leq x_i \leq \pi$	-4.687658
13	Exponential*	40	$-1 \leq x_i \leq 1$	-1
14	Goldstein and Price	2	$-2 \leq x_i \leq 2$	3
15	Griewank*	25	$-600 \leq x_i \leq 600$	0
16	Gulf Research	3	$[0.1, 0, 0] \leq \mathbf{x} \leq [100, 25.6, 5]$	0
17	Hartman 3	3	$0 \leq x_i \leq 1$	-3.862782
18	Hartman 6	6	$0 \leq x_i \leq 1$	-3.322368
19	Helical Valley	3	$-10 \leq x_i \leq 10$	0
20	Hosaki	2	$[0, 0] \leq \mathbf{x} \leq [5, 6]$	-2.345812
21	Kowalik	4	$0 \leq x_i \leq 0.42$	3.074871×10^{-4}
22	Levy and Montalvo 1	3	$-10 \leq x_i \leq 10$	0
23	Levy and Montalvo 2	10	$-5 \leq x_i \leq 5$	0
24	McCormick	2	$[-1.5, -3] \leq \mathbf{x} \leq [4, 3]$	-1.913223
25	Meyer and Roth	3	$-20 \leq x_i \leq 20$	4.355269×10^{-5}
26	Miele and Cantrell	4	$-1 \leq x_i \leq 1$	0
27	Modified Langerman	10	$0 \leq x_i \leq 10$	-0.965
28	Modified Rosenbrock	2	$-5 \leq x_i \leq 5$	0
29	Multi-Gaussian	2	$-2 \leq x_i \leq 2$	-1.296954
30	Neumaier 2	4	$0 \leq x_i \leq 4$	0
31	Neumaier 3*	30	$-900 \leq x_i \leq 900$	-4930
32	Odd Square	20	$-15 \leq x_i \leq 15$	-1
33	Paviani	10	$2.001 \leq x_i \leq 9.999$	-45.77845
34	Periodic	2	$-10 \leq x_i \leq 10$	0.9
35	Powell Quadratic	4	$-10 \leq x_i \leq 10$	0
36	Price Transistor Monitoring	9	$-10 \leq x_i \leq 10$	0
37	Rastrigin*	30	$-5.12 \leq x_i \leq 5.12$	0
38	Rosenbrock*	50	$-30 \leq x_i \leq 30$	0
39	Salomon*	50	$-100 \leq x_i \leq 100$	0
40	Schaffer 1	2	$-100 \leq x_i \leq 100$	0
41	Schaffer 2	2	$-100 \leq x_i \leq 100$	0
42	Schwefel*	40	$-500 \leq x_i \leq 500$	-1.675932×10^4
43	Shekel 10	4	$0 \leq x_i \leq 10$	-10.53628
44	Shekel 5	4	$0 \leq x_i \leq 10$	-10.15320
45	Shekel 7	4	$0 \leq x_i \leq 10$	-10.40282
46	Shekel Foxholes	10	$0 \leq x_i \leq 10$	-10.20879
47	Shubert	2	$-10 \leq x_i \leq 10$	-186.7309
48	Sinusoidal	20	$0 \leq x_i \leq 180$	-3.5
49	Storn Tchebyshev	9	$-256 \leq x_i \leq 256$	0
50	Wood	4	$-10 \leq x_i \leq 10$	0


Table 1. Final set of global optimization test problems, taken from [1]. Problems marked with * had $n = 10$ originally, but had their dimension increased here.

B Further Numerical Results


B.1 Selecting best Py-BOBYQA variant: low accuracy and budget




(a) Comparison of Py-BOBYQA ($p = 2n + 1$) variants on smooth GO test set




(b) Comparison of Py-BOBYQA ($p = \mathcal{O}(n^2)$) variants on smooth GO test set




(c) Comparison of Py-BOBYQA ($p = 2n + 1$) variants on GO test set with multiplicative noise



(d) Comparison of Py-BOBYQA ($p = \mathcal{O}(n^2)$) variants on GO test set with multiplicative noise




(e) Comparison of Py-BOBYQA ($p = 2n + 1$) variants on GO test set with additive noise



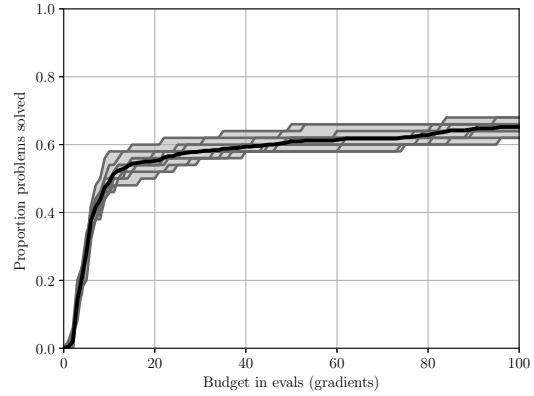
(f) Comparison of Py-BOBYQA ($p = \mathcal{O}(n^2)$) variants on GO test set with additive noise

Figure 9. Comparison of Py-BOBYQA variants on GO test set—low accuracy ($\tau = 10^2$) and budget (10^2 gradients).

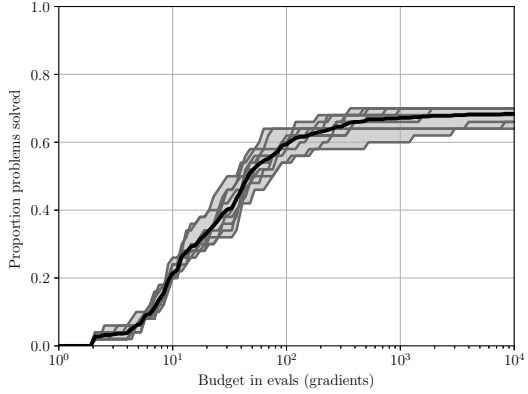
B.2 Py-BOBYQA Variance Testing: $2n + 1$ interpolation points




(a) Fixed restart radius, low accuracy/budget.



(b) Adaptive restart radius, low accuracy/budget.





(c) Fixed restart radius, high accuracy/budget.



(d) Adaptive restart radius, high accuracy/budget.


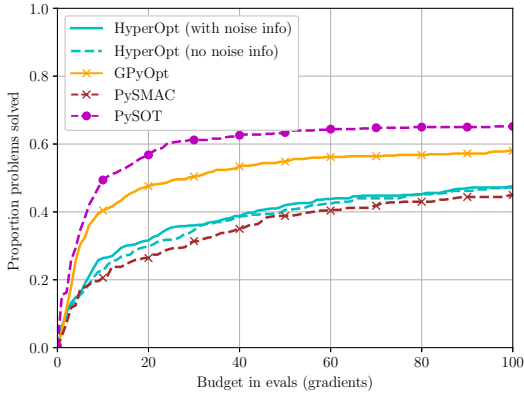
Figure 10. Demonstration of impact of random starting point for Py-BOBYQA ($p = 2n + 1$ with soft restarts); compare with Figure 5 for $p = \mathcal{O}(n^2)$ points. Low accuracy and budget are $\tau = 10^{-2}$ and 10^2 gradients respectively; high accuracy and budget are $\tau = 10^{-5}$ and 10^4 gradients respectively. The dark lines are the average of all runs.

B.3 Selecting the best Bayesian/surrogate solver: noisy objectives



(a) Low accuracy ($\tau = 10^{-2}$), multiplicative noise.

(b) High accuracy ($\tau = 10^{-5}$), multiplicative noise.



(c) Low accuracy ($\tau = 10^{-2}$), additive noise.

(d) High accuracy ($\tau = 10^{-5}$), additive noise.

Figure 11. Comparison of Bayesian and surrogate solvers on the GO test set for noisy problems; see Figure 6 for the smooth case.