# Lawrence Berkeley National Laboratory

**LBL Publications** 

# Title

Designing a Framework for Solving Multiobjective Simulation Optimization Problems

# Permalink

https://escholarship.org/uc/item/3gr340t0

# Authors

Chang, Tyler H Wild, Stefan M

# **Publication Date**

2025

# DOI

10.1287/ijoc.2023.0250

# **Copyright Information**

This work is made available under the terms of a Creative Commons Attribution License, available at <a href="https://creativecommons.org/licenses/by/4.0/">https://creativecommons.org/licenses/by/4.0/</a>

Peer reviewed

# Designing a Framework for Solving Multiobjective Simulation Optimization Problems

Tyler H. Chang<sup>1</sup> and Stefan M. Wild<sup>2</sup>

<sup>1</sup>Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S Cass Ave Bldg 240, Lemont, IL, USA 60439, tchang@anl.gov

<sup>2</sup>Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA, USA 94720, wild@lbl.gov

January 13, 2025

#### Abstract

Multiobjective simulation optimization (MOSO) problems are optimization problems with multiple conflicting objectives, where evaluation of at least one of the objectives depends on a black-box numerical code or real-world experiment, which we refer to as a simulation. While an extensive body of research is dedicated to developing new algorithms and methods for solving these and related problems, it is challenging and time consuming to integrate these techniques into real world production-ready solvers. This is partly due to the diversity and complexity of modern state-of-the-art MOSO algorithms and methods and partly due to the complexity and specificity of many real-world problems and their corresponding computing environments. The complexity of this problem is only compounded when introducing potentially complex and/or domain-specific surrogate modeling techniques, problem formulations, design spaces, and data acquisition functions. This paper carefully surveys the current state-of-the-art in MOSO algorithms, techniques, and solvers; as well as problem types and computational environments where MOSO is commonly applied. We then present several key challenges in the design of a Parallel Multiobjective Simulation Optimization framework (ParMOO) and how they have been addressed. Finally, we provide two case studies demonstrating how customized ParMOO solvers can be quickly built and deployed to solve real-world MOSO problems.

**Keywords:** multiobjective optimization, simulation optimization, engineering design optimization, surrogate modeling, open-source software design

### **1** Introduction and Motivation

Multiobjective optimization problems (MOOPs) are optimization problems involving two or more potentially conflicting objectives. Such problems arise in numerous fields of science, with examples such as multidisciplinary engineering design [1, 2], scientific model calibration [3], high-performance computing (HPC) library autotuning [4], particle accelerator design [5], neural network architecture search [6, 7, 8], and computational chemistry [9, 10]. In such problems the goal is to find a set of "good" points from a design space with respect to a vector-valued objective function  $\mathbf{F} : \mathcal{X} \to \mathbb{R}^o$ . Here,  $\mathcal{X}$  is called the feasible design space and is typically assumed to be a compact simply-bounded subregion of  $\mathbb{R}^n$ . In this paper we assume each objective  $F_i(\mathbf{x})$  is bounded from below so that the feasible objective space  $\mathbf{F}(\mathcal{X})$  is lower bounded in each dimension. In the standard formulation, one seeks to minimize each objective (i.e., component of  $\mathbf{F}$ ), a problem that is written as

$$\min_{\mathbf{x}\in\mathcal{X}}\mathbf{F}(\mathbf{x}).\tag{1}$$

Since it is typically not possible to find one  $\mathbf{x} \in \mathcal{X}$  that simultaneously minimizes all o components of  $\mathbf{F}$ , the solution to (1) is generally a set of design points and corresponding objective values. Rather than

describe when an objective value corresponds to a solution for (1), it is easier to begin by describing when an objective value does *not* correspond to a solution. An objective value  $\mathbf{F}(\mathbf{x}) \in \mathbf{F}(\mathcal{X})$  is said to be dominated if there exists a  $\mathbf{y} \in \mathcal{X}$  such that  $\mathbf{F}(\mathbf{y}) \leq \mathbf{F}(\mathbf{x})$ . Here we use the vector-inequality " $\leq$ " to indicate that  $\mathbf{F}(\mathbf{y})$  is componentwise less than or equal to  $\mathbf{F}(\mathbf{x})$  with strict inequality  $F_i(\mathbf{y}) < F_i(\mathbf{x})$  in at least one component *i*. Conversely, an objective value  $\mathbf{F}(\mathbf{x}^*)$  is said to be nondominated in a set of objective values  $\Omega \subset \mathbb{R}^o$  if for all  $\mathbf{F}(\mathbf{y}) \in \Omega$ ,  $\mathbf{F}(\mathbf{y}) \leq \mathbf{F}(\mathbf{x}^*)$ .

If  $\mathbf{F}(\mathbf{x}^*)$  is nondominated in  $\mathbf{F}(\mathcal{X})$ , then  $\mathbf{F}(\mathbf{x}^*)$  is said to be Pareto optimal for (1), and  $\mathbf{x}^*$  is said to be efficient for (1). The set of all Pareto optimal objective values typically forms a (o-1)-dimensional trade-off surface, called the Pareto front (along a subset of the boundary of  $\mathbf{F}(\mathcal{X})$ ). For further details on general MOOPs, we refer readers to the book by [11]. In general, a design point and objective value pair  $(\mathbf{x}^*, \mathbf{F}(\mathbf{x}^*))$  are in the solution set for (1) if they are efficient/Pareto optimal. In a typical multiobjective optimization application, a *decision maker* uses a multiobjective optimization solver to identify a subset of all Pareto optimal objective values and select one or more of these as the preferred solution(s) based on domain expertise and personal preference.

A large class of real-world MOOPs consists of those derived from expensive numerical simulations [3, 5, 1, 2], computer experiments [4, 6, 7, 8], and real-world experiments [12, 13, 9, 10]. We refer to such problems as multiobjective simulation optimization (MOSO) problems. Often, the simulations in a MOSO problem are black-box processes that are expensive to evaluate and do not admit derivative information. While the objectives themselves may have components that are less expensive to evaluate and/or admit derivative information, the MOSO problem as a whole is typically treated as a black box and solved by using black-box optimization techniques [14]. Furthermore, due to the expected expense of evaluating the simulations, the performance of MOSO algorithms is typically measured in terms of solution accuracy per number of total simulation evaluations. In this paper we focus on designing solvers for MOSO problems. Consequently, here we use the convention that "simulation" refers to an expensive black-box function that yields only zerothorder information. As described later in this paper, solving MOSO problems does not always preclude the usage of some partial derivative information. However, we restrict ourselves to techniques that account for the presence of a computationally expensive black-box process somewhere in the computing chain, and we work under the assumption that the complete gradient for all objectives and constraints will not be available.

From an optimization researcher's perspective, solving MOSO problems efficiently requires coordination between state-of-the-art and emerging numerical techniques from a variety of fields of study, including designof-experiments, surrogate modeling, scalarization, uncertainty quantification, and optimization algorithms. We elaborate on many of these techniques in Section 2.1.

From a practitioner's perspective, solving real-world MOSO problems involves more than numerical algorithms. First, even within the narrow field of MOSO problems, there can be significant variation in problem definitions. This includes application and domain-specific modeling techniques, problem structures, and feasible design spaces. Second, it is important to consider the resources and computing environment where the problem will be solved. This includes wet-lab environments and HPC resources, any of which may require additional steps and technologies to efficiently operate. We elaborate on these challenges in Section 2.2.

Coordinating all of the above techniques from different fields of study into usable solvers that can be deployed to solve real problems is a massive and time-consuming undertaking, which often limits the rate at which novel methods can be adopted. The Parallel Multiobjective Simulation Optimization (ParMOO) library implements a framework based on surrogate modeling for solving MOSO problems, with the goal of providing an easy-to-use and modular interface for implementing and deploying a wide variety of MOSO algorithms and emerging techniques. ParMOO was originally introduced in [15], with a focus on the quality of the open-source software processes, intended user-base, and usability. This paper introduces the considerations that have continued to drive ParMOO's design and development, and the MOSO algorithms and techniques that ParMOO currently supports. We begin by deconstructing surrogate-based MOSO solvers into a system of simpler modules for performing subtasks that are common across many solver architectures. We then present an abstract framework for integrating these modules, while reducing complexities. To demonstrate ParMOO's effectiveness in exploiting real-world problem structures, we conclude by providing access to two computationally inexpensive benchmark problems exhibiting some of the challenging features discussed above. These problems are easy to test on, since the expensive simulations have been replaced by machine learning models trained on real-world datasets.

The remainder of the paper is organized as follows. In Section 2 we review existing techniques, problem types, trends, and software in MOSO, and we describe the technical challenges and trade-offs involved in building a general-purpose MOSO library that achieves all of our design goals. In Section 3 we introduce our framework for addressing these challenges. In Section 4 we perform a brief performance study and validation on a well-known multiobjective optimization test problem. We then present two case studies that demonstrate the flexibility and power of our framework: Section 5 involves the calibration of a physics model, and Section 6 involves the design of material in a wet-lab environment. We conclude in Section 7 with a short discussion of our results.

## 2 Background and Design Principles

We now lay out the primary goals and requirements when designing a general framework for implementing modern surrogate-based MOSO algorithms. These goals and requirements depend on an understanding of the current state-of-the-art techniques and modern trends in MOSO, so we begin by reviewing these.

There are many different families of multiobjective optimization algorithms based on where in the optimization process the decision maker expresses their preference for a solution. In this paper, we are interested in *a posteriori* methods, where the decision maker expresses their preference after solving the problem. Other methods such as *a priori* methods and *interactive* methods require different techniques. They are summarized in the survey by [16].

It is useful to further divide a posteriori MOSO methods into three broad categories that are common in the literature: multiobjective evolutionary algorithms and metaheuristics (MOEAs), which rely on random combinations and/or perturbations to a so-called, "pool," "archive," or "population" of previously evaluated design points; multiobjective direct search (MDS) methods, which rely on (often deterministically) evaluating promising design points around one or more current iterates in search of multiobjective directions; and multiobjective Bayesian optimization (MBO), which focuses on minimizing a multiobjective *acquisition function* to select design points for evaluation. In Section 2.1, we summarize into each of the above methods, and the current representative algorithms for solving MOSO problems. Additionally, there are other methods – mostly based on either generic surrogate modeling or response surface methodology (RSM) – which do not fit cleanly into any of the above categories. In Section 2.1, we explore these techniques as well.

#### 2.1 Fundamental Methods for Solving MOSO Problems

We now summarize the key techniques that are prevalent among a posteriori MOSO algorithms, focusing on those that are most relevant to our work; for a complete survey, see [17].

We begin by discussing scalarization, which is a classical technique that is still prevalent in many modern MOSO algorithms and applications. Scalarization is the basis for almost all a priori methods, as it collapses a MOOP into a single-objective subproblem that can be (approximately) solved with a single-objective optimization solver. In a priori methods, the decision maker may provide some fixed set of scalarization functions, which are carefully selected to target only the desired solution point(s), as pre-specified by the decision maker. However, in the context of a posteriori methods, a family of multiple scalarization functions can be used to produce a discrete approximation to the efficient set and/or Pareto front.

A scalarization is defined by a scalarization function  $Q: \mathbb{R}^o \to \mathbb{R}$ , such that

$$\arg\min_{\mathbf{x}\in\mathcal{X}}Q(\mathbf{F}(\mathbf{x}))\subset\arg\min_{\mathbf{x}\in\mathcal{X}}\mathbf{F}(\mathbf{x}).$$

The scalarized problem  $\min_{\mathbf{x} \in \mathcal{X}} Q(\mathbf{F}(\mathbf{x}))$  can be solved by using any derivative-free optimization method. The most commonly used scalarization technique is the weighted sum scalarization [11, Ch. 3], where the objective is collapsed via a weighted average. Other common scalarization techniques include the epsilonconstraint method [11, Ch. 4], Pascoletti-Serafini scalarization [18], the reference point method [19], quadratic scalarization schemes [20], and (augmented) weighted Chebyshev scalarization [21].

In order to provide an approximation to the complete Pareto front, each of the above-mentioned approaches would need to be combined with some *adaptive scheme* for sweeping through different members in a family of scalarization functions, thus producing numerous solution points covering the Pareto front [22, 23, 18, 24]. Although it is a classical technique, scalarization is still widely used in MOEAs [25, 26],

MDS [27, 23], and MBO methods [28]. Even algorithms that do not explicitly rely upon scalarization still use some of these scalarizations indirectly. For example, the MDS method MOIF [29] and the line search technique DFMO [30] both select a multiobjective steepest descent direction that could be viewed through the lens of scalarization.

Many newer MOSO algorithms (particularly in MOEAs) have replaced the scalarization function with an *indicator function*, which focuses on assessing the quality of the Pareto front approximation when adding a new point  $\mathbf{F}(\mathbf{x})$ . Given a dataset  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{F}(\mathbf{x}_i))\}_{i=1}^N$  of previously evaluated design points and objective values, an *indicator function*  $I(\mathcal{D})$  calculates the quality of the Pareto front approximation given by the nondominated objective values in  $\mathcal{D}$ . Designing a good indicator function is also nontrivial, but it is generally agreed that a good indicator should produce solution points that are on or close to the true Pareto front, offer good coverage of the entire Pareto front, and are well distributed [31, 32]. By choosing a scalarization function  $Q_{\mathcal{D}}(\mathbf{x}) = I(\mathcal{D} \cup \{(\mathbf{x}, \mathbf{F}(\mathbf{x}))\}$  or  $Q_{\mathcal{D}}(\mathbf{x}) = -I(\mathcal{D} \cup \{(\mathbf{x}, \mathbf{F}(\mathbf{x}))\}$ , the indicator can play the part of both the scalarization function and its adaptive scheme. The key difference here is that an indicator-functiondriven scalarization depends upon the dataset  $\mathcal{D}$  in addition to the value of  $\mathbf{F}(\mathbf{x})$ . One of the most commonly used indicators in MOSO algorithms is the hypervolume improvement indicator, which measures the amount of hypervolume between a potential objective value and previously observed objective values [33, 34]. This indicator is used as the basis for many multiobjective algorithms, such as Archived Multiobjective Simulated Annealing (AMOSA) [35].

Since each scalarization function reduces a MOSO problem to a single-objective subproblem, this technique must be coupled with a single-objective solver. Such a solver could implement a single-objective evolutionary algorithm, heuristic, direct search method, Bayesian optimization solver, or any other singleobjective approach. It is also possible to apply a different MOSO solver (or a single iteration thereof) to the MOSO problem, then select the solutions that best minimize the scalarization a posteriori. For a recent survey of derivative-free optimization methods that could be used in the above context, see [36].

For MOSO problems where simulation evaluations are expensive, the total problem expense can get out of hand if too many simulation evaluations are required. Therefore, many modern MOSO solvers reduce the need for true simulation evaluations by combining one or more of the above-mentioned multiobjective algorithms with a computationally cheaper surrogate model [37, 32, 38]. Here, each component of an expensive-to-evaluate objective  $\mathbf{F}(\mathbf{x}) = (F_1(\mathbf{x}), \ldots, F_o(\mathbf{x}))$  is modeled by a computationally cheaper surrogate function  $\hat{\mathbf{F}}_{\mathcal{D}} = (\hat{F}_1, \ldots, \hat{F}_o)$ , based on the current dataset  $\mathcal{D}$ , and such that  $\hat{F}_1 \approx F_1, \ldots, \hat{F}_o \approx F_o$ . Several "candidate solutions" can then be suggested by finding points that are efficient/Pareto optimal for the surrogate problem

$$\min_{\mathbf{x}\in\mathcal{X}} \left( \hat{F}_1(\mathbf{x}), \dots, \hat{F}_o(\mathbf{x}) \right). \tag{2}$$

When  $\mathbf{F}$  is sufficiently more expensive to evaluate than  $\hat{\mathbf{F}}_{\mathcal{D}}$  in terms of computational time and resources, the cost of finding (approximately) efficient/Pareto optimal points for (2) can be much less than the cost of finding such points for (1). In practice, these surrogates could be any of a variety of classical approximation or machine learning [39] models.

The main challenge when using surrogate models is that the surrogate models must be asymptotically accurate for the optimization algorithm to convergence to an accurate solution. Typically, the accuracy of the surrogate depends on the geometry and density of the dataset  $\mathcal{D}$  in  $\mathcal{X}$ , the componentwise smoothness of the underlying function  $\mathbf{F}$ , and the choice of surrogate model [40]. In most applications, the smoothness of  $\mathbf{F}$  is an unknown but fixed quantity. Therefore, we must rely on either the density and geometry of  $\mathcal{D}$  or uncertainty information returned from the surrogate models in order to ensure acceptable surrogate accuracy.

In the context of model-based derivative-free optimization (DFO), one focuses on the local properties of  $\mathcal{D}$  around the current iterate  $\mathbf{x}^{(k)}$ . As more design points are evaluated during the course of the optimization algorithm, the density of  $\mathcal{D}$  rapidly increases near local minima. Therefore, the focus is on maintaining an acceptable conditioning for the surrogate modeling problem in the neighborhood of such minima, which is determined by the local geometry of the design points in  $\mathcal{D}$  [41]. This geometry can be maintained by occasional requiring the optimization algorithm to evaluate model-improving design points instead of strictly model minimizing points [42]. Since these techniques, focus on local quality of the surrogate model, the surrogate problem (2) must be constrained with a local trust region (LTR) so that only candidate solutions in the LTR are suggested by the optimization algorithm. There are currently a few surrogate-driven MOSO

algorithms that exactly follow this approach [43, 44], but none with software to our knowledge. Other existing algorithms use aspects of this approach or are similar in spirit [32, 45]. From the single-objective setting, this "model-based DFO" approach is known to produce fast converging highly scalable algorithms, but only guarantees local convergence [46]. However, this approach can be coupled with global search or multi-start techniques if global accuracy is desired [47].

On the other hand, Bayesian optimization leverages uncertainty models specific to surrogates. Certain surrogates, such as Gaussian processes, admit an uncertainty function  $\Sigma_{\hat{\mathbf{F}},\mathcal{D}} : \mathcal{X} \to \mathbb{R}^o_+$  such that  $\Sigma_{\hat{\mathbf{F}},\mathcal{D}}(\mathbf{x}) \propto |\mathbf{F}(\mathbf{x}) - \hat{\mathbf{F}}_{\mathcal{D}}(\mathbf{x})|$ , (where ' $\propto$ ' can be understood as meaning "componentwise proportional") [48]. In MBO, the scalarization/indicator functions are augmented with the value of  $\Sigma_{\hat{\mathbf{F}},\mathcal{D}}$  to assess the overall utility of evaluating a potential candidate design point  $\mathbf{x}$ , in terms of both improving the Pareto front approximation and maintaining global accuracy of the surrogate model [49, 50]. Notably, the PareGO algorithm uses the uncertainty information from a Gaussian process model to calculate the expected improvement in the augmented Chebyshev scalarization [28], and the (q)EHVI algorithm uses (a quasi-stochastic approximation to) the expected improvement in the hypervolume indicator [51]. MBO produces globally convergent methods, but with the drawback of a slower rate of convergence, particularly in high-dimensional design spaces. Therefore, many methods aimed at high-dimensional design spaces use localization strategies such as LTRs, similarly to the model-based DFO methods above [52].

An alternative framework from the statistics literature is multiobjective RSM. RSM begins with a "search" step, which generates and evaluates a large design-of-experiments [53] or quasi-random sample [54] to explore the design space and collect an initial dataset for surrogate modeling. Once enough data has been collected to ensure sufficiently accurate surrogate modeling, the typical multiobjective RSM approach applies a polynomial surrogate model, multiple scalarizations, and an optimization procedure to solve all scalarized surrogate problems [13, Ch. 7]. Several existing multiobjective surrogate-modeling algorithms use variations or multiple iterations of this approach, including VTMOP [32], SOCEMO [38], and PAWS [45]. Multiobjective RSM can be efficient in many-objective settings since a large portion of RSM's computational cost comes from the initial search step, whose results can be shared across all scalarizations. Therefore, the incremental cost of solving for many objectives by applying many scalarizations is not much greater than the cost of solving for a single scalarization by RSM. The drawback of this approach is that in high-dimensional spaces, it may not be feasible to generate a sufficiently large design-of-experiments to guarantee global surrogate accuracy. Still, all of the above methods also require an initial surrogate-modeling dataset, which is often gathered by similar techniques as in RSM; see, e.g., the search techniques used in the Bayesian optimization library BoTorch [55].

In practice, the lines between the types of algorithms discussed above are often blurred. Furthermore, not all MOSO algorithms use all of the methods described above. In contrast, ParMOO focuses on techniques that utilize:

- An acquisition function or similar technique for setting targets in the objective space;
- An optimization solver (or single iteration thereof) for generating the next iterate based on the acquisition function;
- A surrogate model for approximating expensive functions and either an uncertainty function or model improvement procedure for maintaining its accuracy; and
- A search technique for exploring the design space and generating the initial samples for surrogate modeling.

Prior to the release of ParMOO, there was no open-source software for implementing generic MOSO solvers that adhere to the above structure. [15] focused on the usage of ParMOO rather than the underlying algorithm, framework, and techniques supported. In this paper, we present the design strategies and architecture used for solving MOSO problems via any algorithm that falls into the above framework, with a focus on how such combinations can be most effectively applied in domain-specific applications.

#### 2.1.1 Other Techniques in MOSO.

As previously stated, not all MOSO algorithms utilize all of the techniques discussed in Section 2.1.

One of the biggest departures from Section 2.1 is that not all MOSO algorithms use a technique that can be cast under the umbrella of an acquisition function. However, most require a technique for selecting the point or set of points to iterate from. For example, the MDS method MODIR [56] generalizes the single-objective DIRECT algorithm by way of a multiobjective identification procedure. The MDS methods DMulti-MADS [57], DMS [58], and MultiGLODS [59] use some combination of indicator functions and/or trust-region radii to select each iterate. In the context of MOEAs, a dominance-based sorting metric is typically applied during population selection [60, Ch. 2.3.3]. Such an approach is taken in the well-known MOEA NSGA-II [61]. While many of these selection metrics could be calculated from the information available to an acquisition function (in particular, the dataset  $\mathcal{D}$ ), these methods represent a significant departure from the techniques discussed in Section 2.1. Additionally, it is worth noting that many of the algorithms listed above do not rely on surrogate modeling as a core feature, which is another significant departure from the techniques covered toward the end of Section 2.1.

While not all encompassing, the class of techniques covered in Section 2.1 give the core ingredients for many MOSO algorithms. Algorithms of this nature are the primary focus of this paper.

#### 2.2 Design Challenges when Building MOSO Solvers

We now explore the key design challenges addressed by ParMOO. ParMOO is designed to use the techniques laid out in Section 2.1, specifically, acquisition/indicator/scalarization functions, surrogate modeling, and design-of-experiments, to solve a diverse set of scientific MOSO applications. Therefore, we focus on challenges related to these techniques and how they can be deployed in real-world applications and scientific computing environments.

#### 2.2.1 State-of-the-Art and Domain-Specific Techniques.

First, as highlighted in Section 2.1, state-of-the-art surrogate-driven MOSO solvers require integrating several different techniques, including acquisition functions, optimization solvers, surrogate models, and design-of-experiments. This challenge is compounded when we consider the inclusion of domain- and problem-specific techniques such as machine learning surrogates with physics-informed constraints [62] or acquisition functions designed to target extremely rare events [63].

In ParMOO, we seek to support most such techniques, including domain-specific methods that we may not yet be aware of. This requires a clean application programmer's interface (API) for users to implement new, domain-specific methods. This, in turn, requires a standardization of APIs for each individual method and clear definition of how these methods will interact, which may not be easy to define for all of the techniques in Section 2.1.

#### 2.2.2 Structured Problem Formulations.

Beyond using domain- and problem-specific algorithms, one can go a step farther and explicitly model domain- and problem-specific structures in the MOSO problem formulation. This includes modeling composite problem definitions, as described in [64, 65, 66], where the objective can be expressed as the composition of one or more functions. In particular, we can modify the MOOP problem definition from (1), using

$$\min_{\mathbf{x} \in \mathcal{X}} \mathbf{F}(\mathbf{x}, \mathbf{S}(\mathbf{x})) \tag{3}$$

where  $\mathbf{F} : \mathcal{X} \times \mathbb{R}^m \to \mathbb{R}^o$  and  $\mathbf{S} : \mathcal{X} \to \mathbb{R}^m$ . In this formulation, it is assumed that  $\mathbf{F}$  is an *algebraic* function that is cheap to evaluate and whose gradient function may be known, while  $\mathbf{S}$  is a computationally expensive simulation function, as previously discussed. This formulation decouples the "simulation" aspect of the computation from the objective function, allowing for the usage of structured non black-box optimization solvers.

One common example for  $\mathbf{F}$  would be a sum-of-squares function  $F_i(\mathbf{x}, \mathbf{S}(\mathbf{x})) = \sum_{j \in J} S_j(\mathbf{x})^2$ , for one or more component functions  $F_i$  and an index set J. In the single-objective case, this structure can be exploited by Gauss-Newton inspired methods to achieve super-linear convergence rates [67], as is the case in the derivative-free least-squares solver POUNDERS [68]. While the theoretical advantage may not be as strong, explicitly modeling the composition in (3) typically leads to convergence with fewer simulation evaluations for a variety of other algebraic **F**-functions [64, 65, 66].

Another structure of  $\mathbf{F}$  that is specific to the multiobjective setting is a heterogeneous problem definition [43], where one or more  $F_i$  is an algebraic function with no dependence on an expensive simulation output, but other  $F_i$  are either the direct outputs or calculated from a simulation output. This sort of problem definition has been acknowledged in previous works [32, 43], but we are not aware of a formal analysis. The typical approach in the context of surrogate-based MOSO, is to apply surrogate models to the components that depend upon simulation outputs and use the algebraic formulations instead whenever they are available. This approach lends well to the abstraction in (3).

In ParMOO, we support problem definitions of the form (3), which is a slight departure from the standard black-box optimization methodology. However, it is important to note that this formulation still involves a black-box process, and therefore cannot be solved with purely gradient-based methods.

#### 2.2.3 Changes to the Design Space – Mixed Variables and Constraints.

Thus far, we have assumed that  $\mathcal{X}$  is a compact, simply-bounded subregion of  $\mathbb{R}^n$ . However, widely used MOSO software often support broad categories of design variables and constraints, including various combinations of real (continuous-valued), discrete (integer-valued), and/or categorical (nonordinal) variables, as well as linear and nonlinear constraints [55, 69, 70, 71, 6, 72].

For mixed-type design variables, it is important to note that changing the domain  $\mathcal{X}$  can result in a completely different classification of problem requiring different techniques. For example, when  $\mathcal{X}$  consists of exclusively binary decision variables, a multiobjective ranking and selection algorithm would generally be used [73], which significantly differs from the continuous MOSO algorithms described in Section 2.1. However, many MOSO problems that are largely continuous have a small number of discrete variables, which is something that can be reasonably supported without significantly changing the kind of techniques used.

Since we are focused on methods that were designed for continuous MOSO, the primary challenge is in handling the discrete design variables. For MOEAs, this would be supported by providing a mixedvariable mating procedure to any existing MOEA [71]. For MBO and many MDS techniques it could be sufficient to define a distance, kernel, or neighborhood function that supports integer and categorical variables [74]. For a generic surrogate-based MOSO algorithm, however, one needs a function  $E: x_{i'} \to [0,1]^{\ell}$  for embedding the discrete design variable  $x_{i'}$  into a continuous latent space  $[0,1]^{\ell}$ , where  $\ell$  is the dimension of the embedding. Perhaps the most common example of a latent space embedding is a one-hot encoding, which maps each category to either a 0 or 1, and assigns a value 1 if and only if the corresponding  $x_{i'}$ represents the corresponding category. This latent space is then modeled by the surrogate, which can be used to relax the embedded design variables to a continuous representation. This approach has been widely used in existing surrogate modeling software [39, 75]. For any of the above techniques, it can be difficult to scale generic techniques for problems with a large number of categorical design variables. Therefore, many domain-specific software packages and algorithms utilize domain-specific design space embedder or distance functions [6, 10]. As an example in the context of material design, the Bayesian optimization solver EDBO [10] uses the molecular descriptor calculator Mordred [76] to calculate a three-dimensional latent space encoding molecule networks, and then solves a Bayesian optimization problem in this latent space.

Similarly, the introduction of linear and nonlinear constraints requires us to relax the previous assumption that  $\mathcal{X}$  is simply bounded. In practice, such constraints could be many combinations of known (e.g., part of the problem formulation) or hidden (e.g., discovered at runtime), algebraic or dependent upon simulation outputs, relaxable in the sense that the simulation can still be evaluated when the constraint is violated, and quantifiable in the sense that we can quantify how much it has been violated or not [77]. Clearly, we will not be able to support all combinations of the above, however, we seek to support as many combinations as reasonably possible.

Given the composite objective function definition in (3), it is natural and simplifying to mirror this with a composite constraint definition  $\mathbf{G}(\mathbf{x}, \mathbf{S}(\mathbf{x})) \leq \mathbf{0}$ , where " $\leq$ " denotes componentwise less than or equal to. In this sense, we will support both simulation-based and algebraic inequality constraints and immediately preclude hidden constraints. Then, we can decompose the feasible region  $\mathcal{X}$  into a simply-bounded region  $\mathcal{B}$  (with *n* dimensions and possibly mixed variable types), where the bound constraints are unrelaxable, and

the relaxable nonlinear constraints defined by G. This produces ParMOO's final problem formulation

$$\min_{\mathbf{x}\in\mathcal{B}} \mathbf{F}(\mathbf{x}, \mathbf{S}(\mathbf{x})) \text{ subject to } \mathbf{G}(\mathbf{x}, \mathbf{S}(\mathbf{x})) \leq \mathbf{0}.$$
(4)

For generic nonlinear black-box constraints, there have been several approaches in the single-objective optimization literature for handling the nonlinear constraints  $\mathbf{G}$ , including barrier functions [78] and augmented Lagrangian approaches [79]. Both of these approaches can be thought of as producing a penalty function describing the constraint violation, which can be added to the objective to encourage feasibility. However, the methods for achieving this differ. To extend these methods to the multiobjective case, it is sufficient to apply a cumulative penalty function to all objectives [80]. Considering the options, barrier functions tend be more flexible to apply without significant modification to the optimization algorithm. Among the barrier methods, the two common methods are an extreme barrier, which adds an infinite penalty any time a constraint is violated, and a progressive barrier, which adds a penalty based on the constraint violation and progressively increases this penalty over time.

In ParMOO, we seek to support a limited number of mixed variables and relaxable nonlinear constraints **G**. We pursue this through embeddings and barrier functions, which requires limited modification to the techniques outlined in Section 2.1.

#### 2.2.4 Parallel and Other Computing Environments.

Another challenge in MOSO is achieving efficient resource utilization in diverse computing environments. Many existing optimization solvers provide parallelism through a single paradigm (e.g., Python multiprocessing, OpenMP, or MPI) [71, 32, 72]. However, this may not be sufficient when dealing with diverse computing environments including distributed systems, HPCs, laboratory clusters, and even wet-lab environments, any one of which may not support one or more common parallel computing paradigms.

In order to address this complexity, many platforms [81, 12] and libraries/frameworks [82, 83] have been created for coordinating parallel evaluations, complex simulations, data storage, and heterogeneous resource utilization. Additionally, domain-specific tools for managing simulation, modeling, and optimization interplay also manage distribution over parallel resources [84]. However, effectively utilizing these platforms is nontrivial if the optimization software was not originally designed with mindfulness for these paradigms [85, 86].

In ParMOO, we have designed with parallelism and simulation environment flexibility in mind, being sure to decouple simulation evaluation from the solvers and techniques.

#### 2.2.5 Maintainability and Usability.

A recent movement in the general scientific computing space is to reduce technical debt and improve scientist productivity through better scientific software development practices [87]. Many scientific software libraries are adopting such practices [82, 46], and several existing multiobjective optimization libraries have recently been refactored to improve maintainability [88, 69].

One of our key goals in the design of ParMOO has been to build a software package that provides an API for implementing the techniques discussed in the previous two sections, without sacrificing maintainability and usability.

#### 2.2.6 Design Goals.

In summary to the challenges and requirements laid out in this section, we now list five key goals in the design of ParMOO.

- 1. We want to provide a highly customizable software framework for building and deploying surrogatebased MOSO solvers.
- 2. We want to give the optimization solver access to any available structure in how the simulation outputs are used to define the objectives.
- 3. We want to be flexible in our support for a wide variety of problem types, including mixed variables and constraints.

- 4. We want to make our framework easy to deploy in a wide variety of scientific workflows.
- 5. Our software framework and workflow must be easy to use, maintain, and extend.

Goals 1, 2, and 3 require careful thought in terms of what kinds of MOSO algorithms, problem types, and structures are supported and how these techniques will interact. Goals 4 and 5 are engineering problems, which constrain the breadth of methods that we can reasonably support when addressing Goals 1, 2, and 3. Although it constrains our decision making on the techniques we can use, ParMOO's workflow and general project structure (some of the main components of Goal 5) have previously been extensively discussed in [15]. Therefore, we limit our discussion regarding Goal 5 to where it affects our decision making for the other design goals.

#### 2.3 Existing MOSO Software

In addition to the software laid out below, we acknowledge the existence of other special-purpose solvers for domain-specific variations of the MOSO problem. In this work, however, we do not consider solvers that are tied to a specific application, such as neural architecture search [8] and chemical experiment design [10], or solvers that target other variations of the problem (1), including online multiobjective optimization [89] and multiobjective reinforcement learning [90]. We also acknowledge that a similarity exists between surrogatebased MOSO and active learning [91], which we consider to be a more generic problem than surrogate-based MOSO.

In this section, we exclusively list open-source, production-quality software packages for solving MOSO problems. The term "production quality" here focuses on software implementations that are advanced enough for a non-expert to use in a non-academic application. Hallmarks of such software include some combination of detailed user or API documentation, publication in a software journal, and a large user base. Although often not distinguished, it is also important to clarify that not every well-known MOSO algorithm has a one-to-one mapping with such a software implementation.

Much of the available software implements MOEAs. Widely used MOEA-based solvers include ParEGO [28], which also integrates Gaussian process surrogate modeling, and SPEA2 [92]. There are also many libraries, including Platypus [93], pymoo [71], PlatEMO [94], jMetal/jMetalPy [95, 69], and pagmo/pygmo [70]. The majority of these libraries implement all of the well-known MOEAs, such as NSGA-II [61] and NSGA-III [25, 26]. Although not specific to MOOPs, the framework DEAP [96] is also frequently used for implementing distributed evolutionary algorithms in Python, including distributed MOEAs.

Packages with MDS methods include the following. The DFO-lib [97] contains serial Fortran implementations of MODIR and DFMO and both Matlab and Python implementations of MOIF. Serial and parallel Fortran solvers are distributed in VTMOP [32], which utilizes RSM together with adaptive weightings and trust-region methods. NOMAD is a parallel-capable C++ library of industrial-grade single- and multiobjective direct search methods, which until recently only contained the biobjective solver BiMADS [72]. In a newer release of NOMAD v4 [88], an implementation of DMulti-MADS was added as well. BoostDFO [98] is a MATLAB library containing parallel implementations of single- and multiobjective direct search solvers, including DMS and MultiGLODS. PyMOSO [99] is a Python framework that is targeted primarily at integer-valued problems.

In recent years several Bayesian optimization libraries and frameworks have emerged, supporting MBO. The Python library Dragonfly [6], was designed for solving neural architecture search problems but can be applied to generic MBO problems. Perhaps most relevant to our goals in this paper is the BoTorch [55] framework for implementing and deploying parallel, production-ready, generic Bayesian optimization (including multiobjective) algorithms while layering over the automatic differentiation PyTorch framework [100].

To summarize this section, we present a list of the solvers and libraries discussed thus far in Table 1. This table represents the relevance of these libraries in relationship to ParMOO's design goals. Therefore, we focus on the features offered by these libraries that relate to Design Goals 1–4 from Section 2.2.6. Notably, support for domain- and application-specific problem definitions (related to Goal 2) is not shown since, of the methods listed, only BoTorch and ParMOO support this kind of problem formulation at this time. We acknowledge that many of these software packages are under active development and may add features in the future, so we base our classifications by the documented features at the time of this publication.

Name	Type	Language	Method	Constr.	Var. Types	Surrogates
NOMAD v4	L	C++	MDS	yes	mixed	yes
BoostDFO	L	Matlab	MDS	some	real	yes
BoTorch	L	Python	MBO	yes	mixed	yes
DEAP	L	Python	MOEA	yes	mixed	no
DESDEO	L	Python	any	yes	real	yes
Dragonfly	L	Python	MBO	yes	mixed	yes
jMetal/jMetalPy	L	Java/Python	MOEA	yes	mixed	no
DFO-lib	L	Fortran/Python/Matlab	MDS	some	real or int	no
ParEGO	S	$\mathbf{C}$	MOEA/MBO	no	real	yes
pagmo	L	C++	MOEA	some	mixed	no
ParMOO	L	Python	MDS/MBO	yes	mixed	yes
PlatEMO	L	Matlab	MOEA	some	mixed	some
Platypus	L	Python	MOEA	yes	mixed	no
pygmo	L	Python	MOEA	some	mixed	no
pymoo	L	Python	MOEA	some	mixed	no
PyMOSO	L	Python	MDS	yes	$\operatorname{int}$	no
SPEA2	S	$\mathbf{C}$	MOEA	no	real	no
VTMOP		Fortran	MDS	no	real	yes

Table 1: General-purpose open-source MOSO software. For each software package, the columns are labeled as follows: "Type" indicates whether this is an individual solver (S) or library of solvers (L); "Language" specifies the primary development language; "Method" classifies each software as primarily using MOEA, MDS, or MBO; "Constr." indicates whether nonlinear constraints are supported; "Var. Types" indicates the types of variables supported; and "Surrogates" indicates whether surrogate modeling is used.

# 3 A Framework for Multiobjective Simulation Optimization

We now outline our framework for solving MOSO problems. In particular, we describe how this framework addresses the design goals and challenges laid out in Section 2.2.6. The framework is implemented in the software library ParMOO [15], which can be obtained by following the installation instructions in its online documentation [101].

### 3.1 An Object-Oriented Modular Design

Despite the level of complexity necessary to achieve Goal 1, the biggest challenge of this work is maintaining a usable, maintainable, and extensible API while managing the complexity of these interacting techniques (Goal 5). To do so, we focus on a modular design, where each module addresses a particular component of the larger problem through a common interface. The natural paradigm for implementing such a framework is object-oriented programming (OOP). To this end, we have implemented abstract base classes (ABCs) for each of the solver components. Referring back to Section 2.1, the four primary ABCs for ParMOO are

- the GlobalSearch class, which abstracts the API for utilizing design-of-experiments and other sampling techniques;
- the SurrogateFunction class, which abstracts the API for fitting, updating, and evaluating a generic surrogate function, plus either evaluating the uncertainty function or suggesting model improvement points, as applicable;
- the AcquisitionFunction class, which abstracts the API for implementing an acquisition function (including scalarization and indicator functions) and selecting the next iterate (if applicable); and
- the SurrogateOptimizer class, which accepts a scalarized, differentiable or non-differentiable function and uses an appropriate method to produce the next iterate or (if applicable) call for a model improvement.

To allow users to utilize ParMOO without implementing each of these techniques, we provide a built-in library of implementations for several standard techniques from each of the categories.

At the center of all of this, the MOOP class stores a potentially complex MOSO problem definition of the form (4) and coordinates implementations of the ABCs to solve a MOSO problem. This is achieved using the "builder" design pattern [102]. The MOOP class itself is also highly modular and extensible, with clearly documented methods for performing both private and public tasks. Developers can easily overwrite certain functionalities, such as the method that distributes expensive simulation evaluations, in order to adapt to novel architectures and use cases. In this way, the modular OOP framework also addresses goal 5.

The downside of this approach is that the object-oriented API adds an additional level of complexity, similar to that of learning a new modeling language, such as JuMP [103] or Pyomo [104]. It also requires users to become aware of the basic components of a surrogate-based multiobjective solver, in order to correctly utilize our library of components and framework. Ultimately, this mental overhead is significantly greater than that required to learn to properly utilize other Python libraries of multiobjective solvers, such as pymoo [71] and pagmo/pygmo [70]. However, we believe that this approach is a necessary compromise in order to achieve Goals 1–3 as laid out in Section 2.2.6.

In this way, ParMOO is not comparable to many of the libraries and individual solvers listed in Table 1. Of these software packages, ParMOO is perhaps most comparable to DEAP [96] and BoTorch [55], which implement similar frameworks for solving slightly different classes of problems using different classes of techniques.

#### 3.2 Distinguishing between Simulations, Objectives, and Constraints

The next challenge that we need to address directly is Goal 2, which requires making simulation outputs and function evaluations visible to the optimization solver, for the problems defined in (3) and (4). Our key insight here is to manage and model simulation outputs separately from objective values and constraint penalties. Figure 1 shows how ParMOO treats MOOPs, with a black-box simulation function **S** mapping into an intermediate "simulation output space" S, before **F** maps  $\mathcal{X} \times S$  into the feasible objective space.



Figure 1: The feasible design space  $\mathcal{X}$  (left) is mapped into an intermediate simulation output space  $\mathcal{S}$  (center) via the black-box simulation **S**, then to the feasible objective space  $\mathcal{Y}$  (right) via the algebraic objective function **F**.

The abstraction depicted in Figure 1 allows us to handle the problem formulations in (3) and (4). Instead of applying surrogate modeling to the objectives using a dataset  $\mathcal{D}$  of design point, objective value pairs, (e.g., fitting  $\hat{\mathbf{F}}_{\mathcal{D}}(\mathbf{x}_i) \approx \mathbf{F}(\mathbf{x}_i)$ ) we model surrogates of the *intermediate simulation space*  $\hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x}_i) \approx \mathbf{S}(\mathbf{x}_i)$ using a dataset  $\mathcal{D}$  of design point, *simulation output* pairs. Then, similarly as in (2), we use a solver to approximately minimize an acquisition function or scalarization of the surrogate problem

$$\min_{\mathbf{x}\in\mathcal{B}} A_{\mathcal{D}}(\mathbf{F}(\mathbf{x}, \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})), \boldsymbol{\Sigma}_{\hat{\mathbf{S}}, \mathcal{D}}(\mathbf{x})) \text{ subject to } \mathbf{G}(\mathbf{x}, \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})) \leq \mathbf{0}$$
(5)

where  $\Sigma_{\hat{\mathbf{S}},\mathcal{D}}$  is the uncertainty function for  $\hat{\mathbf{S}}_{\mathcal{D}}$  given  $\mathcal{D}$ .

In each iteration, the SurrogateOptimizer object has access to the scalarized outputs  $A_{\mathcal{D}}(\mathbf{F}(\mathbf{x}, \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})), \boldsymbol{\Sigma}_{\hat{\mathbf{S}}, \mathcal{D}}(\mathbf{x}))$ , the objective function  $\mathbf{F}$ , and the simulation surrogate  $\hat{\mathbf{S}}_{\mathcal{D}}$ , as well as the constraint violation function  $\mathbf{G}$ . This information is sufficient to implement a wide variety of composite-structure-exploiting optimization solvers.

### 3.3 Managing Complexity through Problem Embeddings

To address Goal 3, we need to handle problems involving continuous, integer-valued, and categorical variables, as well as nonlinear constraints, without requiring any problem-specific techniques or heavy modification to the structure outlined in Section 3.1. To maintain usability and maintainability (goal 4), we must be careful in the techniques that we use.

As discussed in Section 2.2.3, for handling mixed variables, our method of choice is to embed the problem into a continuous latent space and then solve a relaxation of the problem. ParMOO abstracts this idea by creating a hidden embedding layer based on built-in or user-provided embedder/extractor functions  $\mathbf{E}_{in}$  and  $\mathbf{E}_{out}$  such that  $\mathbf{E}_{in}(\mathbf{x}) : \mathcal{X} \to [0,1]^{\ell}$  and  $\mathbf{E}_{out}(\mathbf{E}_{in}(\mathbf{x})) = \mathbf{x}$ ; see Figure 2. Upon input, all design variables are passed through their respective embedders,  $\mathbf{E}_{in}$ , and any variables taking on discrete values are relaxed via the surrogate model. After ParMOO solves the surrogate optimization problem, all candidate design points are extracted from the latent space back into the original design space via  $\mathbf{E}_{out}$  and binned to their nearest legal values if necessary. ParMOO provides a default embedding for categorical variables, which utilizes a one-hot encoding, followed by a dimension reduction procedure to eliminate unused dimensions of the latent space (i.e., two different latent variables representing the same categorical variable cannot both have nonzero values at the same time). However, this technique is not scalable to a large number of categorical variables. Therefore, for problems with large numbers of categorical variables, the default embedding is based upon a simple rescaling to the range [0, 1].



Figure 2: A depiction of the data flow in ParMOO. The embedder module takes possibly discrete values from the design space and maps them to a normalized latent space  $[0,1]^{\ell}$ . The extractor module accepts candidate designs in the range  $[0,1]^{\ell}$  from the optimization module, and maps them back into  $\mathcal{X}$ .

The primary advantage of this approach is that it maintains a level of simplicity in the API and allows us to focus on continuous optimization techniques regardless of the problem type. In particular, there is no need to use customized search, surrogate modeling, or optimization techniques for solving mixed-variable problems (although doing so may improve performance) since the default techniques are designed to operate on the rescaled latent space. The flexibility to accept a custom embedding tool is also useful for domains such as molecular engineering, where these tools are available [76]. As previously discussed, for problems involving large numbers of categorical variables without any domain-specific embedding tools, the default method may become inefficient and it may be advantageous to take a completely different approach that is not based on continuous MOSO algorithms.

In order to incorporate a *p*-dimensional nonlinear constraint function **G**, we apply a cumulative penalty function,  $P(\mathbf{x}) = \lambda \sum_{i=1}^{p} \max(G_i(\mathbf{x}, \mathbf{S}(\mathbf{x})), 0)$  or  $P(\mathbf{x}) = \lambda \sum_{i=1}^{p} \max(G_i(\mathbf{x}, \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})), 0)$ , to the surrogate problem. This penalty is added to *all* of the objective functions when solving (5). Combining with the embedder/extractor functions described above, we can reduce a nonlinearly constrained, partially discrete feasible design space  $\mathcal{X}$  to a bound-constrained, well-conditioned, continuous MOSO problem

$$\min_{\mathbf{x}\in[0,1]^{\ell}} A_{\mathcal{D}}(\mathbf{F}(\mathbf{E}_{\text{out}}(\mathbf{x}), \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})) + \lambda \sum \max(\mathbf{G}(\mathbf{E}_{\text{out}}(\mathbf{x}), \hat{\mathbf{S}}_{\mathcal{D}}(\mathbf{x})), \mathbf{0}), \boldsymbol{\Sigma}_{\hat{\mathbf{S}}, \mathcal{D}}(\mathbf{x})),$$
(6)

where the max is taken elementwise. The penalty parameter  $\lambda > 0$  is progressively increased, on an exponential schedule, whenever the solution to (6) is not expected to be feasible. This is similar to the progressive barrier approach described in [78], and has been successfully applied with simulation-based constraints in [5]. This approach has been shown to be effective when the constraints are hard to satisfy [105].

One of the downsides of this approach, is that the penalty function uses constraint violations, which implicitly requires the constraint functions in  $\mathbf{G}$  to be both relaxable and quantifiable. Therefore, in ParMOO, only the upper/lower bound constraints on each design variables can be unrelaxable, which is a necessary compromise to achieve reasonable performance on heavily constrained problems.

#### **3.4** Flexibility and Extensibility to Novel Workflows

Goal 4 was to make our framework flexible enough that it is easy to deploy in a variety of scientific workflows. Key examples include HPC and wet-lab environments.

To achieve this goal, we have designed ParMOO to issue simulation evaluations exclusively via the wrapper function MOOP.evaluateSimulation(), which is exclusively called from within the solver routine MOOP.solve(). In situations where only the simulation command must be changed for integration with existing workflow technology, this can be achieved by extending the MOOP class and overwriting the MOOP.evaluateSimulation() method. In other situations, where control over the frequency with which simulation evaluations are distributed is required (for example, when batching simulation evaluations), the entire MOOP.solve() method can be overwritten. The latter approach is taken in ParMOO's libE\_MOOP class, which extends the MOOP class and overwrites the solve method to dynamically distribute simulation evaluations on HPC systems using the libEnsemble library [82, 83]. The libE\_MOOP class is currently the recommended method for achieving scalable parallelism with ParMOO.

#### 3.5 The Resulting Framework

Putting everything together, we present our framework for solving MOSO problems as implemented in ParMOO. In each iteration, ParMOO uses Algorithm 1 to generate the next batch of candidates (implemented in the MOOP.iterate() method).

To fill in the various components of Algorithm 1, the user must create and fully populate a MOOP object. When initializing the MOOP, the user must specify a SurrogateOptimizer for suggesting candidates for solving or iterating upon (6) and a dictionary of hyperparameters. Next, the user must add problem details, including

- *n* design variable dictionaries, each specifying a design variable for the problem and (if applicable) a custom Embedder and Extractor routine;
- s simulation dictionaries, each defining a simulation/data source for the problem, plus the number of intermediate outputs  $(m_i)$ , the GlobalSearch technique used to sample that simulation's output space, and the SurrogateFunction used to model that simulation's output and either evaluate uncertainties or perform model improvement iterates;
- *o* objective dictionaries, each specifying an algebraic function that can be used to calculate the objective values from the design variables and simulation outputs and (optionally) the gradient of that objective;
- p constraint dictionaries, identical to the objective dictionaries but for the purpose of evaluating relaxable constraint violations (i.e., the **G** functions); and
- q acquisition function dictionaries, each specifying an AcquisitionFunction that can be used to scalarize the problem, where q also determines the batch size for parallel evaluations.

A UML diagram outlining this framework is shown in Figure 3. Note that only the relevant public methods discussed in this section are shown. In the MOOP class's implementation, several additional public and private helper methods are included, as well as additional "setter" methods to facilitate saving, loading, logging, and checkpointing [101].

Algorithm 1: A single iteration of ParMOO's MOOP.iterate() method

- **input** : k > 0 is the current iterate.
- **input** :  $q_0 \ge 0$  is an initial search budget.
- **input** : q is the batch size, i.e., the number of AcquisitionFunctions.
- **input** : n is the dimension of the design space and  $\ell$  is the dimension of the latent space.
- input :  $\mathbf{E}_{in} : \mathcal{X} \to [0,1]^{\ell}$  is the cumulative embedder function for all design variables.
- input :  $\mathbf{E}_{out} : [0,1]^{\ell} \to \mathcal{X}$  is the extractor function (inverse of  $\mathbf{E}_{in}$ ).
- input  $: \mathbf{S} : \mathcal{X} \to \mathbb{R}^m$  is the simulation function.
- input :  $\mathbf{F} : \mathcal{X} \times \mathbb{R}^m \to \mathbb{R}^o$  is the objective function.
- input  $: \mathbf{G} : \mathcal{X} \times \mathbb{R}^m \to \mathbb{R}^p$  is the constraint function.
- **input** :  $\mathcal{D}^{(k)}$  is the set of all  $(\mathbf{x}, \mathbf{S}(\mathbf{x}))$  pairs evaluated prior to iteration k.
- input : GlobalSearch.search( $q_0, \ell$ ) is a procedure that generates a design-of-experiments or sample of size  $q_0$  in  $[0,1]^{\ell}$ .
- input : SurrogateFunction.fit( $\mathcal{D}$ ) fits the surrogates  $\hat{\mathbf{S}}_{\mathcal{D}} : [0,1]^{\ell} \to \mathbb{R}^m$ , and (if applicable) their uncertainty function  $\Sigma_{\hat{\mathbf{S}},\mathcal{D}} : [0,1]^{\ell} \to \mathbb{R}^m$ .
- input : SurrogateFunction.improve( $\mathcal{T}$ ) produces a model improving step for  $\hat{\mathbf{S}}_{\mathcal{D}}$  in  $\mathcal{T} \subset [0,1]^{\ell}$ (if applicable).
- input :  $\{A_{\mathcal{D}}^{(i)}\}_{i=1}^{q}$  is a set of q AcquisitionFunctions.
- input : SurrogateOptimizer.solve $(A_{\mathcal{D}}, \mathbf{F}, \mathbf{G}, \hat{\mathbf{S}}_{\mathcal{D}}, \boldsymbol{\Sigma}_{\hat{\mathbf{S}}, \mathcal{D}}, \mathbf{x}^{(0)})$  solves (6), iterates toward the solution to (6) starting from  $\mathbf{x}^{(0)}$ , or calls the SurrogateFunction.improve( $\mathcal{T}$ ) method and provides a LTR  $\mathcal{T}$ .

**output:** The batch  $\mathcal{C}^{(k)}$  of candidate design points after iteration k.

```
1 \mathcal{C}^{(k)} \leftarrow \emptyset;
```

```
2 if k = 0 then
        \mathcal{C}' \leftarrow \texttt{GlobalSearch.search}(q_0, \ell);
3
```

```
for
each \mathbf{y} \in \mathcal{C}' do
4
     5
6
 else
```

```
\mathcal{D}' \leftarrow \{ (\mathbf{E}_{in}(\mathbf{x}), \mathbf{S}(\mathbf{x})) : \text{ for all } (\mathbf{x}, \mathbf{S}(\mathbf{x})) \in \mathcal{D}^{(k)} \};
7
```

```
\hat{\mathbf{S}}_{\mathcal{D}}, \boldsymbol{\Sigma}_{\hat{\mathbf{S}},\mathcal{D}} \leftarrow \mathtt{SurrogateFunction.fit}(\mathcal{D}');
8
```

```
for i \leftarrow 1 to q do
  9
                            \mathbf{x}^{(0,i)} \leftarrow \arg\min_{(\mathbf{x},\mathbf{S}(\mathbf{x}))\in\mathcal{D}'} A_{\mathcal{D}'}(\mathbf{F}(\mathbf{x},\mathbf{S}(\mathbf{x})),\mathbf{0}));
10
```

```
\mathbf{y}^{(k,i)} \leftarrow \texttt{SurrogateOptimizer.solve}(A_{\mathcal{D}'}, \mathbf{F} \circ \mathbf{E}_{\text{in}}, \mathbf{G} \circ \mathbf{E}_{\text{in}}, \hat{\mathbf{S}}_{\mathcal{D}}, \boldsymbol{\Sigma}_{\hat{\mathbf{S}},\mathcal{D}}, \mathbf{x}^{(0,i)});
```

```
if \mathbf{y}^{(k,i)} \notin \mathcal{C}^{(k)} then
```

```
\mathcal{C}^{(k)} \leftarrow \mathcal{C}^{(k)} \cup \{\mathbf{E}_{\text{out}}(\mathbf{y}^{(k,i)})\};
```

```
else
```

```
14
                              \mathbf{y}' \leftarrow \texttt{SurrogateFunction.improve}([0,1]^{\ell});
15
                              \mathcal{C}^{(k)} \leftarrow \mathcal{C}^{(k)} \cup \{\mathbf{E}_{\text{out}}(\mathbf{y}')\};
16
```

```
17 return \mathcal{C}^{(k)};
```

11

12

13

#### 3.6Limitations of the Framework

We have endeavored to cover a great number of MOSO algorithms and techniques and make them easily accessible for a wide variety of applications. While we have succeeded for a wide breadth of algorithms, our framework is not without limitations. For MOSO algorithms that do not use acquisition functions or surrogates, as discussed in Section 2.1.1, it would likely be inefficient to implement these in ParMOO's framework and thus shoehorn their iterations into Algorithm 1. Additionally, as acknowledged in Section 2.2.3, ParMOO's framework is fundamentally based upon techniques for continuous, bound-constrained MOSO and handles mixed variables and nonlinear constraints by reducing these problems to bound-constrained,



Figure 3: UML diagram outlining the key dictionaries, components, and methods that make up a MOOP object and its contents.

continuous MOSO problems. For problems with no continuous variables or unrelaxable constraints, this is likely not the most effective approach and other frameworks and classes of algorithms may be better suited.

Although MOSO is typically a derivative-free endeavor, we have made efforts to offer limited support for partial derivative information, since the algebraic functions  $\mathbf{F}$  and  $\mathbf{G}$  could be differentiable and ParMOO's optimization solvers can use their gradient information when available. However, this can become misleading in the sense that ParMOO is designed to solve problems involving at least *some* computationally expensive black-box simulations. If there is no black-box simulation function involved in the problem definition and if gradients are available for all components of the objective and constraint definitions, then other approaches not covered here would likely be more effective.

Finally, it is worth noting that ParMOO is fundamentally a framework, not an algorithm. Therefore, much of the work in exploiting problem structures and implementing quality algorithms is left to the user. While ParMOO's framework and development is mature, the variety of techniques available is still under active development and not all of the techniques listed in Section 2.1 are available at this time. Careful analysis and clever algorithms are needed to gain a theoretical advantage when exploiting most problem structures, although we will show in Section 5 and Section 6 that ParMOO can still gain some clear advantage in most applications with relatively little work.

### 4 Performance Summary

In this section we apply a trust-region solver built in ParMOO to a common academic test problem from the literature in order to assess its parallel scaling. Before doing so, we would like to caution about the limitations of this study. First, as described in Section 3, ParMOO is a framework and software library for implementing solvers, but not an algorithm in and of itself; there can be significant variations in ParMOO's performance depending on the methods used. Second, the focus of this paper and ParMOO's design is on achieving performance on real MOSO applications. Since the academic problem used here does not necessarily reflect a real-world MOSO application, performance on this problem should be understood as validating that ParMOO achieves reasonable performance on solvable problems, and not that the particular ParMOO solver described is "better" or "worse" than existing methods. Thus, we have intentionally chosen a single relatively easy problem from the multiobjective optimization literature, the DTLZ2 problem from the widely-used DTLZ test suite [106]. We use the variation of this problem with 3 objectives and 10 design variables. The problem has no constraints, no mixed variables, and no composite structures of the form shown in Figure 1 to exploit. The simplicity of the problem allows us to focus on the correctness of our implementation and parallel scaling under generic conditions.

Since the entire DTLZ2 problem is algebraic, we must create an artificial simulation. Therefore, we have created a wrapper for the DTLZ2 function that turns it into a black-box simulation by creating an artificially long runtime. This is done by uniformly sampling a runtime between one and three seconds, then waiting for that amount of time using Python's built-in time.sleep() function. The objectives (**F**-functions) are then identity mappings from this artificial simulation's three output fields.

As our sample solver, we have created a parallel ParMOO solver for this problem by instantiating the libE\_MOOP builder class. The solver has been populated using the TR\_LBFGSB implementation of the SurrogateOptimizer to solve a trust-region-constrained surrogate problem with L-BFGS-B; a combination of the randomized epsilon-constraint scalarization RandomConstraint and fixed-weight weighted sum scalarization FixedWeight for our AcquisitionFunctions; a GaussRBF (Gaussian RBF) SurrogateFunction; and a Latin hypercube GlobalSearch. By varying the number of AcquisitionFunction objects, we are able to create variations of this ParMOO solver that will generate batch sizes of 8, 16, and 32 candidate simulations per iteration. Since ParMOO parallelizes simulation evaluations, this batch size determines the largest amount of parallelism available in a single iteration.

Since the Latin hypercube sampling and default model-improvement procedures are randomized, we have performed 5 runs of the above-defined ParMOO solver with different random seeds for each run. We have solved with each random seed and each batch size serially and with 2-, 4-, and 8-way parallelism during simulation evaluation with a budget of 1,000 total simulation evaluations.

To validate our results, we have calculated the widely-used hypervolume performance indicator at the end of each run [31], averaged these values, and compared across runs. For the hypervolume indicator, larger values are better, with the maximal value being the total volume between the Pareto and a reference point (in this case, the point (1, 1, 1)). Hypervolume values are notoriously difficult to interpret, so for comparison, we have also solved DTLZ2 on the same budget using the widely-used NSGA-II implementation in pymoo [61, 71].

The results of the comparison are shown in Table 2. Even without exploiting any structures, this shows that ParMOO solver's performance on DTLZ2 (according to the hypervolume indicator) is consistently higher than pymoo after 1,000 simulation evaluations, indicating that ParMOO achieves at least reasonable performance on academic benchmark problems. There is slight variation between the performance of the ParMOO solvers of different batch sizes, however, we would recommend against drawing conclusions on such limited data, especially since DTLZ2 is not a particularly representative application for ParMOO.

Method	pymoo	ParMOO-8	ParMOO-16	ParMOO-32
Hypervolume	0.28	0.33	0.33	0.37

\_

Table 2: Hypervolume indicator for pymoo and ParMOO with batch sizes 8, 16, 32. Larger values are better.

ParMOO does not change Algorithm 1 when running in parallel, it only parallelizes the simulation evaluations in a batch. Therefore, we have not observed any changes in its convergence (as measured by the hypervolume indicator) when increasing the number of threads other than minor and uncorrelated fluctuations due to random number generation. Therefore, if we are able to reduce the walltime to perform a 1,000 simulation run with any of the above solvers, this marks a distinct parallel advantage. In Figure 4, we consider the total walltimes when running each ParMOO batch size with increasing number of threads. Note that since the average simulation time is 2 seconds, the expected total walltime when performing 1,000 total simulation evaluations serially would be approximately 2,000 seconds plus any iteration costs incurred by ParMOO.



Figure 4: Walltimes when performing 1,000 1-3 second simulation evaluations in ParMOO (with batch sizes 8, 16, and 32) with increasing number of threads.

As seen in Figure 4, ParMOO iteration computations incur a bit of overhead for this solver configuration and problem size (note that the single threaded walltime is a bit over 2,000 seconds) but the walltime to solution decreases proportionally with the number of threads. This parallelism is slightly less than perfect strong scaling and is limited by the candidate batch size. Still, for most reasonable solver configurations, when performing 8 simulations in parallel requiring an average 2 seconds per simulation ( $\pm 1$  second variation), one could expect to perform 1,000 simulations in ParMOO in just under 500 seconds.

### 5 Case Study: Solving a Multiobjective Inverse Problem

In this section we apply ParMOO to a multiobjective inverse problem, where the goal is to tune the parameters of the Fayans energy density functional (EDF) model based on experimental data. This problem is fully described in [3]. The forward model is expensive to evaluate and not publicly available. Therefore, to maintain reproducibility of results, we optimize a synthetic problem based on a neural network model that was trained on the dataset above.

#### 5.1 Background on Fayans EDF Calibration

We now review important aspects of the Fayans EDF calibration. Let  $T_{\mathbf{x}} : s \to t$  denote the Fayans EDF model. A single forward evaluation of  $T_{\mathbf{x}}$  accepts a normalized parameter vector  $\mathbf{x} \in \mathbb{R}^{13}$  and an input s and produces a physical observation t. Given 198 observational data pairs  $(s_1, t_1), \ldots, (s_{198}, t_{198})$ , the goal of the Fayans EDF model calibration is to find "good" parameter vectors  $\mathbf{x}^* \in \mathbb{R}^{13}$  such that  $T_{\mathbf{x}^*}(s_i) \approx t_i$  relative to a given standard error  $\sigma_i > 0$  for  $i = 1, \ldots, 198$ .

In traditional Fayans EDF calibration, it is assumed that all observations have been normalized by the standard errors  $\sigma$  such that they have equal and independent normalized errors. It is then reasonable to minimize the  $\chi^2$  loss across all 198 observations via the single-objective formulation

$$\min_{\mathbf{x}\in\mathbb{R}^{13}}\sum_{i=1}^{198} \left(\frac{T_{\mathbf{x}}(s_i) - t_i}{\sigma_i}\right)^2.$$
(7)

In [3], however, it is noted that the contributions to the  $\chi^2$  loss may vary across 9 different observational types. Since the 9-objective formulation is prohibitively expensive to solve, we consider a simplified three-objective formulation, where related observational types are combined, leaving just three observational classes. This allows us to define the following multiobjective formulation of the Fayans EDF calibration:

$$\min_{\mathbf{x} \in \mathbb{P}^{13}} F_{\mathbf{x},j}, \qquad j = 1, 2, 3,\tag{8}$$

where  $F_{\mathbf{x},j} = \sum_{i \in \Phi_j} \left( \frac{T_{\mathbf{x}}(s_i) - t_i}{\sigma_i} \right)^2$  and  $\Phi_1, \Phi_2, \Phi_3$  is a partitioning of  $\{1, \ldots, 198\}$  based on the three observational classes described above.

In this section we solve the multiobjective formulation of the Fayans EDF calibration problem (8). In this problem, evaluation of the forward model for all 198 observations  $[T_{\mathbf{x}}(s_1), \ldots, T_{\mathbf{x}}(s_{198})]$  is viewed as a single computationally expensive black box. The three loss functions  $F_{\mathbf{x},j}$  (j = 1, 2, 3) all have a composite structure, specifically, a sum-of-squares structure.

#### 5.2 The Neural Network Residual Model

For our experiments, since  $T_{\mathbf{x}}$  is expensive and we strive for reproducibility, we fit a multilayer perceptron (MLP) to approximate the standardized residual functions  $R_i(\mathbf{x}) = \frac{T_{\mathbf{x}}(s_i)-t_i}{\sigma_i}$ ,  $i = 1, \ldots, 198$ . We do this using a dataset of labeled observations for the residual functions  $R_1(\mathbf{x}), \ldots, R_{198}(\mathbf{x})$  for 52,079 distinct values of  $\mathbf{x}$ . This dataset was gathered by [3] and, as described in that paper, comes from running multiple starting design points with several different single-objective solvers on the single-objective formulation of the problem (7). Therefore, we note that the dataset is not uniformly distributed and has increased density in the neighborhood of several local solutions to the single-objective problem (7).

The primary challenge when training the MLP used in this section is ensuring that the prediction accuracy is high for optimal values of the multiobjective problem (8). In particular, large regions of the parameter space were determined to be unstable for the single-objective formulation, and therefore no observational data is available. Additionally, within the stable region, for several values of i, the observational data for  $R_i(\mathbf{x})$  could range in magnitude from less than 1 to over  $10^{20}$ . While prediction errors that are greater than 1 in magnitude would be acceptable (and expected) for these large residuals, it is essential for our fidelity to the original problem that we maintain prediction errors on the order of  $10^{-2}$  in the neighborhood of the true solutions.

To handle infeasible regions of the parameter space where no observational data is available, we impose bound constraints for ParMOO, following the constraints given by [3], shown in Table 3. To handle the wide range on observational values, we apply a double-logarithmic transformation to large values of  $R_i(\mathbf{x})$ , followed by a tanh transformation. Because this collapses huge error values, this transformation ensures that, during training, the MLP places much higher importance on matching observations where the residual function  $R_i(\mathbf{x})$  has a low score.

For the network architecture, we define an MLP with 13 inputs, 198 outputs, and 2 hidden layers, with 256 nodes per layer. We then apply tanh activation functions for every layer of the network (including the output layer) and train the network using the transformed data defined above. For validation, we first stratify the complete observational database by low/high residuals across all three observational classes defined in (8) and then withhold 5% of the data across all stratifications for our validation set. The network was trained in keras [107] using 5,000 epochs of RMS-prop. We verify that the trained model obtains low relative error across all stratified residual ranges in the validation set. Most importantly, after descaling the outputs to their original ranges, on the bottom stratification where the residuals are lowest in magnitude across all three observational classes, we obtain a mean absolute error (MAE) of just 0.036 on the validation set, which is acceptable accuracy for this problem.

In the remainder of this section we use the trained keras model described above as a synthetic representation of the original Fayans EDF calibration problem.

#### 5.3 Exploiting Structure in ParMOO

The sum-of-squares structure is well studied in the single-objective black-box optimization literature. In typical software implementations of this strategy [68], special care is taken to ensure that all geometric

Variable Name	Lower Bound	Upper Bound
$ ho_{eq}$	0.146	0.167
E/A	-16.21	-15.50
K	137.2	234.4
J	19.5	37.0
L	2.20	69.6
$h_{2-}^v$	0.0	100.0
$a^s_+$	0.418	0.706
$h^{s}_{ abla}$	0.0	0.516
$\kappa$	0.076	0.216
$\kappa'$	-0.892	0.982
$f_{ex}^{\xi}$	-4.62	-4.38
$h^{\xi}_+$	3.94	4.27
$h_{ abla}^{\dot{\xi}}$	-0.96	3.66

Table 3: Upper and lower bounds on the stable region for the parameter space for the Fayans EDF model  $T_{\mathbf{x}}(s)$ , as reported by [3]. We have access to observational data only in these ranges, and therefore the trained MLP will be accurate only within these bounds.

conditions on the interpolation nodes hold, thus requiring the algorithm to occasionally take geometryimproving evaluations and using a sequence of local (trust-region) approximations. However, it has also been shown that simply modeling this composite structure typically leads to fewer simulation evaluations in practice [64, 65, 66].

In ParMOO's built-in LocalGaussRBF surrogate model implementation, the trust-region radius is adaptively chosen based on the distance to the (n + 1)th nearest neighbor of the current iterate. Then, whenever the solution to the trust-region-constrained surrogate problem fails to produce a sufficient decrease in the scalarized objective value, the model-improving step is triggered, which draws a random sample from a distribution whose variance is highest in directions of low variance in the current simulation dataset. This approach does not guarantee any asymptotic convergence rate advantage (indeed, such a convergence rate would be difficult to even define in the multiobjective setting). We will see empirically, however, that simply modeling  $\mathbf{R}(\mathbf{x})$  separately from  $\mathbf{F}(\mathbf{x})$  greatly accelerates the practical convergence of our structure-exploiting multiobjective solver.

#### 5.4 The ParMOO Fayans EDF Solvers

To demonstrate its ability to utilize parallel function evaluations (as we would need to if we were evaluating the true Fayans EDF model instead of our synthetic MLP model), we use ParMOO's libEnsemble interface to distribute simulation evaluations with a batch size of 10. This is achieved by using 9 epsilon-constraint acquisition functions that target the solutions to (8) and one fixed weighting that targets the solution to the single-objective formulation in (7), since it is equivalent to the  $\chi^2$  loss. To define the problem, we create a libE\_MOOP object in ParMOO and add 13 continuous design variables, with names and bounds as given in Table 3.

To define a structure-exploiting variant as described in the preceding section, we then add a single simulation function that models all 198 components of our synthetic MLP residual function using LTR constrained Gaussian radial basis function (RBF) models and initializes its observational database using a 2,000-point Latin hypercube sample (LHS). We next add three differentiable algebraic objective functions, which calculate the sum of squares from the simulation outputs across each of the three different observational classes. Then, to allow ParMOO to focus on interesting regions of the design space where none of the observational classes result in terrible performance, we add three nonlinear constraints, each enforcing that the squared residuals in a particular observable class should not exceed  $10 \times$  its optimal values reported in [3]. ParMOO then solves the surrogate optimization problems within the LTR using the L-BFGS-B [108] implementation in SciPy [109].

To compare with the performance when the structure is not exploited, as would be the case with other existing multiobjective software, we define a "black-box" variation of this problem. For the black-box variant, the 198-output simulation described above was replaced by a 3-output "black-box" simulation, where the sum of squares has already been computed across all three observational classes. Then, for each objective, we provide an identity map from each simulation output to each objective. Otherwise, we define the problem equivalently as with the structured variant.

In both cases, ParMOO is then run for 800 iterations. Since we begin with a 2,000-point LHS search and provide a batch size of 10 points per iteration, this results in a total budget of 10,000 simulation evaluations for each solver. We note that for a typical simulation optimization problem, this budget of 10,000 simulation evaluations is unrealistically large. In fact, because of the complexity of our Gaussian RBF surrogate models and cost of solving the surrogate optimization problem, for this large a budget our iteration costs are extremely high. However, it is worth running to such a large simulation budget in order to understand the solver's performance in the limit.

Note that because of the complexity of solving the surrogate optimization problem and the nonnegligible cost of evaluating the **keras** model, such a solve requires a substantial amount of compute time. To account for randomness in the LHS searches and weight initializations, we perform 5 runs of each solver and average the performance.

For reproducibility, this entire experimental setup, including the trained MLP that was used as a synthetic problem representation, is available at https://github.com/parmoo/parmoo-solver-farm/tree/main/fayans-model-calibration-2022. Note that for ease of use and compatibility reasons, we have transferred our keras model's weights into an equivalent torch model [100], and this is reflected in the GitHub repository.

### 5.5 Results

The results of solving (8) based on the keras/torch residual model are presented here. In order to account for variability resulting from randomness in the initial design-of-experiments, all performance results have been averaged across five random seeds. First, borrowing the metric used in [3], we present the  $\chi^2$  loss across all observable classes in Figure 5. Note that for  $\chi^2$  loss, small values are better. Next, to estimate the multiobjective performance of our methods, we present a rescaling of the hypervolume performance indicator in Figure 6. For the hypervolume indicator, large values are better.

Note that the hypervolume indicator is extremely sensitive to problem scaling and the choice of Nadir point, so its raw values are difficult to interpret. In an effort to normalize values, we present the improvement in hypervolume over that of the initial 2,000-point LHS design, as a percentage of the hypervolume dominated by that original design. Even after this normalization, however, the absolute value of the hypervolume improvement is still difficult to interpret since it is still influenced by the total hypervolume between the true Pareto front and Nadir point. In Figure 6 we see less than a 0.7% increase in total hypervolume. This is because our Nadir point is determined by the lower-bound constraints that ParMOO enforced on the range of interesting values, which were intentionally set to be overly pessimistic. However, the hypervolume improvements of the two methods relative to each other can still be taken as an indicator of relative performance.

For both the  $\chi^2$  loss and hypervolume improvement, the structured solver converges considerably faster than does the black-box solver. With respect to the  $\chi^2$  loss, in just 200 iterations (4,000 simulation evaluations) the structured solver achieves better performance than the black-box solver will in all 800 iterations (10,000 simulation evaluations). This result is to be expected: since one of our acquisition functions specifically targets this solution with fixed scalarization weights, ParMOO behaves similarly to how a structureexploiting single-objective solver would for this performance metric.

With respect to the percent hypervolume improvement (a true multiobjective performance metric), the ParMOO's structured solver also achieves significantly improved performance. However, the convergence appears to be slightly slower than for the  $\chi^2$  loss. This is to be expected since solving the full multiobjective problem is considerably harder than solving a single scalarization. However, especially for limited computational budgets, the performance of the structured solver is still dramatically improved.



Figure 5: Iteration vs.  $\chi^2$  loss when solving the Fayans EDF calibration with ParMOO, exploiting the sumof-squares structure (structured) and with a standard (black-box) approach. The total simulations used by the end of iteration k are calculated as 2000 + 10k.



Figure 6: Iteration vs. percentage hypervolume improvement when solving the Fayans EDF calibration with ParMOO, exploiting the sum-of-squares structure (structured) and with a standard (black-box) approach. The total simulations used by the end of iteration k are calculated as 2000 + 10k.

## 6 Case Study: A Multiobjective Chemical Design

To demonstrate ParMOO's effectiveness with heterogeneous objectives and flexibility to support diverse scientific workflows, in this section we apply ParMOO to solve a materials engineering problem. Here, two of the three objectives are derived from a continuous-flow chemistry experiment that is conducted in a wet-lab environment, while the third objective is algebraic. As in Section 5, to make our results open and reproducible, in this section we use a nonlinear model of the chemical response surface, based on real-world

data collected using nuclear magnetic resonance (NMR) spectroscopy on the solutions from a continuous-flow reactor (CFR).

### 6.1 Background on CFR Material Design

In material manufacturing applications, our goal is to propose a technique for chemical manufacturing that can be used to produce a desired material with high purity at scale. In particular, in this application we are manufacturing the electrolyte 2,2,2-trifluoroethyl methyl carbonate (TFMC) by mixing one of two predetermined potential solvents with one of two predetermined potential bases. Note that for confidentiality reasons, we do not include the true names of these potential solvents and bases in this section; they are simply labeled as S1, S2 and B1, B2, respectively. This does not affect the reproducibility of our results since they are also labeled as such in the nonlinear model that we used to represent the problem in this section.

Our goal is to find an optimal pairing of solvents and bases and conditions (such as reaction time, equivalence ratio, and temperature) for producing a pure solution of TFMC in a CFR. However, in addition to producing a pure solution, we want to be able to use short reaction times, so that we can produce large quantities of TFMC at scale. We expect that this will require us to use higher temperatures, which could activate a side reaction and produce an unwanted byproduct, thereby reducing the purity.

#### 6.2 The Chemical Response Surface Model

In the context of ParMOO, this problem has three continuous design variables and two categorical design variables. These variable names and their ranges/potential values are given in Table 4. The response values of interest are the integrals of TFMC and byproduct production values over a fixed-length time window, as measured by using NMR spectroscopy.

Variable Name	Var. Type	Lower Bound	Upper Bound	Legal Vals.
temperature (T)	continuous	$35 \mathrm{C}$	$150 \mathrm{C}$	N/A
reaction time (RT)	$\operatorname{continuous}$	$60  \sec$	$300  \sec$	N/A
equiv. ratio (EQR)	$\operatorname{continuous}$	0.8	1.5	N/A
solvent	categorical	N/A	N/A	S1,S2
base	categorical	N/A	N/A	B1,B2

Table 4: Variable types, bounds, and legal values for the CFR material optimization problem discussed in this section. Note that continuous variables have bounds, while categorical variables have legal values.

Since we cannot provide access to the CFR for our experiments, in this section we use a nonlinear response surface representing the physical continuous-flow chemistry experiment. These models were fit by using realworld data that was collected by providing a physical CFR/NMR setup as a simulation that ParMOO could query in closed loop. To do so, ParMOO's extensible API was layered over the MDML tool [12], which uses an Apache kafka backend to distribute requests for experiment evaluations to the CFR and collects and returns NMR results directly to an online database that ParMOO can query for simulation results. For more information on how this data was collected, see [110], which describes the collection of a smaller dataset by using an identical experimental setup.

After a budget of 62 experiments, ParMOO had converged on several approximate solutions to the realworld problem, and it was no longer economically viable to continue the real-world experiment since the cost of real-world materials is nonnegligible. The resulting experimental database of design point/integral-value pairs was used in this section to fit the nonlinear models described above. Since this dataset is relatively sparse in the 5-dimensional input space, special care was taken to ensure that the resulting model does not exhibit non-physical behaviors. In particular, to verify our model, we have ensured that both of our response surfaces

- approximate the underlying data with low MAE, particularly for near-optimal design points;
- do not take on negative values anywhere in the feasible design space, which would be physically impossible since the outputs represent time integrals of material production;

- do not have a sum that exceeds the total amount of solvent and base provided for any inputs in the design space; and
- do not take unexpected maxima/minima along the boundaries of the design space, which could be an artifact of a lack of data in those regions, allowing our model to overextrapolate.

To achieve these criteria, we hand-crafted a small number of physically meaningful nonlinear terms for our response surface based on the expected chemistry of the reaction. We then fit the coefficients of these terms to each of the two integrals, by using a combination of generalized linear regression with scikit-learn [111]. The resulting MAE was found to be within the acceptable ranges. Using hyperparameter tuning with the Powell solver from scipy.optimize.minimize [109], we were also able to guarantee that the individual global minima for each model was nonnegative and the sum of the global maxima was within the acceptable range. We then verified that the individual minima and maxima for each model were located in acceptable regions of the design space, which agree with our physical intuition and empirical data.

In the remainder of this section we use these trained response surface models as the true chemical response surfaces for both the TFMC and byproduct integrals.

### 6.3 Heterogeneous Problem Structure

In this problems, two of the three objectives are the result of a true black-box experiment, which must be carried out in a wet-lab environment using real materials and with an extremely restrictive budget. However, the third objective represents the reaction time, which is one of the directly controllable inputs. Therefore, ParMOO is able to directly control this output, and is able to exploit this ability to accelerate its practical convergence.

### 6.4 The ParMOO CFR Material Design Solver

To define the problem, we provide ParMOO with the five design variables shown from Table 4, with their respective variable types and bounds/values. To handle the two categorical variables, we use ParMOO's default categorical variable embedder, which embeds the four distinct combinations of categorical variables into a three-dimensional continuous latent space. Combined with the three continuous design variables, this results in a six-dimensional effective optimization space.

For the structured variation of the problem, ParMOO is provided with the pretrained nonlinear model of the chemical response surface, described in Section 6.2. ParMOO is configured to treat this as a single simulation with two outputs, using a Latin hypercube sample (LHS) with 50 evaluations to produce the initial database and Gaussian RBFs for surrogate modeling. Next, ParMOO is given three objectives, two of which are identity mappings from the simulation outputs and the third of which is an identity mapping from the reaction time design variable. Then, two epsilon-constraint functions and one fixed-weighting are added, resulting in a batch size of three simulation evaluations per iteration.

For comparison and to demonstrate the advantage in exploiting the heterogeneous problem structure, we also provide an identical implementation of ParMOO for this problem, where the third "reaction time" objective is provided to ParMOO as a third black-box simulation output. This results in ParMOO modeling the third output as a black-box function and ignoring the heterogeneous structure, as would occur when using most other off-the-shelf multiobjective black-box optimization solvers. All other settings are identical as in the structured variation defined above.

One of the critical challenges of performing automatic experimentation in the context of material manufacturing is the cost of raw materials and performing real-world experiments. Although we do not have these costs for our computational model of the material response surface, we have run ParMOO with a restrictive budget of just a 50-point initial design followed by 30 iterations with 3 acquisition functions (140 total experiments). By exploiting the heterogeneous problem structure, we hope to still achieve good performance (especially for our cheap objective) with this limited budget.

The response surface model described above and the code for reproducing the experiments presented here are given at https://github.com/parmoo/parmoo-solver-farm/tree/main/cfr-material-design-2022. The results are presented in the next section.

### 6.5 Results

The results of tuning manufacturing conditions based on our chemical response surface are presented here. Again, all performance results have been averaged across 5 unique random seeds.

First, to assess our performance with respect to our one cheap objective, we present the average minimum observed reaction times (in seconds) subject to a 75% purity constraint in Figure 7. Note that for these reaction times, small values are better. Next, as in Section 5.5, we present the improvement in hypervolume as a proportion of the gap between initial hypervolume and total possible hypervolume (with respect to the ideal point) in Figure 8. Again, for the hypervolume improvement, large values are better.



Figure 7: Iteration vs. minimum reaction time that achieves at least a 75% conversion rate when solving the CFR chemical manufacturing problem with ParMOO, exploiting the heterogeneous structure (structured) and with a standard (black-box) approach. The total simulations used by the end of iteration k are calculated as 50 + 3k.

As in Section 5.5, the structured solver greatly outperforms the black-box approach by both metrics. We note that although the time required to achieve over 75% conversion is a physically meaningful convergence metric and demonstrative of ParMOO's ability to directly control the reaction time (the algebraic objective), this time we did not provide any fixed scalarization that would explicitly target this solution, reducing to a single-objective problem.

We note that ParMOO's structure-exploiting solver achieves excellent performance on both this problem and the Fayans problem from Section 5.5 with little additional work from the user, even though these structures are considerably different.

## 7 Discussion and Continued Work

In this paper we have described the design principles behind the design of the MOSO library ParMOO. To summarize, our five main design goals are

- 1. customizability of solvers;
- 2. exploitation of composite structures in MOSO problem formulation;
- 3. flexibility in support for a wide variety of design spaces;
- 4. ease of use in scientific workflows; and



Figure 8: Iteration vs. percentage hypervolume improvement when solving the CFR chemical manufacturing problem with ParMOO, exploiting the heterogeneous structure (structured) and with a standard (black-box) approach. The total simulations used by the end of iteration k are calculated as 50 + 3k.

5. usability, extensibility, and maintainability as an open-source software package.

We have achieved these goals through an object-oriented design that is highly modularized, utilizes an intermediate simulation output space, and embeds complex problems into a continuous latent input space. This framework has been demonstrated on open-source models of two real-world problems. These models and the code for reproducing our results have been shared through our parmoo-solver-farm repository, and a snapshot of all related code at the time of publication is available in [112].

Although these two problems exhibit completely different structures, ParMOO is able to effectively exploit the structure in both cases and to considerably improve the convergence relative to that of a blackbox approach that does not exploit known structure. Notably, defining such a structure-exploiting solver requires little additional work or customization from the user.

## ACKNOWLEDGMENTS

We are grateful to Jared O'Neal, Witold Nazarewicz, and Paul-Gerhardt Reinhard for providing the Fayans functional data and to Joseph Libera, Jakob Elias, Santanu Chaudhuri, and Trevor Dzwiniel for their respective roles in the collection of the CFR materials data. We are also grateful to Jeff Larson, John-Luke Navarro, and Steve Hudson for their advice on best practices in scientific software development and support in our usage of libEnsemble. We are grateful to four anonymous reviewers for their comments, which have improved the presentation of this work.

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research's SciDAC program under Contract Nos. DE-AC02-05CH11231 and DE-AC02-06CH11357.

### References

J. Sobieszczanski-Sobieski, A. Morris, M. Van Tooren, Multidisciplinary Design Optimization Supported by Knowledge Based Engineering, John Wiley & Sons, Ltd., Chichester, UK, 2015.

- [2] W. Zhao, R. K. Kapania, Multiobjective optimization of composite flying-wings with SpaRibs and multiple control surfaces, in: Proc. 2018 Multidisciplinary Analysis and Optimization Conference, AIAA, 2018, p. 3424. doi:10.2514/6.2018-3424.
- [3] R. Bollapragada, M. Menickelly, W. Nazarewicz, J. O'Neal, P.-G. Reinhard, S. M. Wild, Optimization and supervised machine learning methods for fitting numerical physics models without derivatives, Journal of Physics G: Nuclear and Particle Physics 48 (2020) 024001. doi:10.1088/1361-6471/abd009.
- [4] T. H. Chang, J. Larson, L. T. Watson, Multiobjective optimization of the variability of the highperformance LINPACK solver, in: Proc. 2020 Winter Simulation Conference (WSC 2020), IEEE, 2020, pp. 3081–3092. doi:10.1109/WSC48552.2020.9383875.
- [5] N. Neveu, T. H. Chang, P. Franz, S. Hudson, J. Larson, Comparison of multiobjective optimization methods for the LCLS-II photoinjector, Computer Physics Communication 283 (2023). doi:10.1016/ j.cpc.2022.108566.
- [6] K. Kandasamy, K. R. Vysyaraju, W. Neiswanger, B. Paria, C. R. Collins, J. Schneider, B. Poczos, E. P. Xing, Tuning hyperparameters without grad students: Scalable and robust Bayesian optimisation with Dragonfly, Journal of Machine Learning Research 21 (2020) 1-27. URL: http://jmlr.org/papers/v21/18-223.html.
- [7] F. Karl, T. Pielok, J. Moosbauer, F. Pfisterer, S. Coors, M. Binder, L. Schneider, J. Thomas, J. Richter, M. Lang, E. C. Garrido-Merchán, J. Branke, B. Bischl, Multi-Objective Hyperparameter Optimization in Machine Learning – An Overview, Technical Report 4, 2023. doi:10.1145/3610536.
- [8] M. Parsa, J. P. Mitchell, C. D. Schuman, R. M. Patton, T. E. Potok, K. Roy, Bayesian multi-objective hyperparameter optimization for accurate, fast, and efficient neural network accelerator design, Frontiers in Neuroscience 14 (2020) Article No. 667. doi:10.3389/fnins.2020.00667.
- [9] A. M. Schweidtmann, A. D. Clayton, N. Holmes, E. Bradford, R. A. Bourne, A. A. Lapkin, Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives, Chemical Engineering Journal 352 (2018) 277–282. doi:10.1016/j.cej.2018.07.031.
- [10] B. J. Shields, J. Stevens, J. Li, M. Parasram, F. Damani, J. I. M. Alvarado, J. M. Janey, R. P. Adams, A. G. Doyle, Bayesian reaction optimization as a tool for chemical synthesis, Nature 590 (2021) 89–96. doi:10.1038/s41586-021-03213-y.
- [11] M. Ehrgott, Multicriteria Optimization, Lecture Notes in Economics and Mathematical Systems Series, 2 ed., Springer Verlag, Heidelberg, Germany, 2005. doi:10.1007/3-540-27659-9.
- [12] J. R. Elias, R. Chard, J. A. Libera, I. T. Foster, S. Chaudhuri, The manufacturing data and machine learning platform: Enabling real-time monitoring and control of scientific experiments via IoT, 2020 IEEE 6th World Forum on Internet of Things (WF-IoT) (2020) 1–2. doi:10.1109/WF-IoT48130.2020. 9221078.
- [13] R. H. Myers, D. C. Montgomery, C. M. Anderson-Cook, Response Surface Methodology: Process and Design Optimization Using Designed Experiments, 4 ed., John Wiley & Sons, Inc., Hoboken, NJ, USA, 2016.
- [14] C. Audet, W. Hare, Derivative-free and blackbox optimization, Springer Series in Operations Research and Financial Engineering, Springer International, Charm, Switzerland, 2017. doi:10.1007/ 978-3-319-68913-5.
- [15] T. H. Chang, S. M. Wild, ParMOO: A Python library for parallel multiobjective simulation optimization, Journal of Open Source Software 8 (2023) 4468. doi:10.21105/joss.04468.
- [16] T. R. Marler, J. S. Arora, Survey of multi-objective optimization methods for engineering, Structural and Multidisciplinary Optimization 26 (2004) 369–395. doi:10.1007/s00158-003-0368-6.

- [17] S. R. Hunter, E. A. Applegate, V. Arora, B. Chong, An introduction to multiobjective simulation optimization, ACM Transactions on Modeling and Computer Simulation 29 (2019) 1–36. doi:10. 1145/3299872.
- [18] G. Eichfelder, Scalarizations for adaptively solving multi-objective optimization problems, Computational Optimization and Applications 44 (2009) 249–273. doi:10.1007/s10589-007-9155-4.
- [19] A. P. Wierzbicki, Reference point approaches, in: T. Gal, T. J. Stewart, T. Hanne (Eds.), Multicriteria Decision Making: Advances in MCDM Models, Algorithms, Theory, and Applications, Springer US, Boston, MA, 1999, pp. 237–275. doi:10.1007/978-1-4615-5025-9\_9.
- [20] B. Dandurand, M. M. Wiecek, Quadratic scalarization for decomposed multiobjective optimization, OR Spectrum 38 (2016) 1071–1096. doi:10.1007/s00291-016-0453-z.
- [21] R. E. Steuer, E.-U. Choo, An interactive weighted tchebycheff procedure for multiple objective programming, Mathematical programming 26 (1983) 326–344. doi:10.1007/BF02591870.
- [22] I. Das, J. E. Dennis, Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems, SIAM Journal on Optimization 8 (1998) 631–657. doi:10.1137/S1052623496307510.
- [23] S. Deshpande, L. T. Watson, R. A. Canfield, Multiobjective optimization using an adaptive weighting scheme, Optimization Methods and Software 31 (2016) 110–133. doi:10.1080/10556788.2015. 1048861.
- [24] M. Laumanns, L. Thiele, E. Zitzler, An efficient, adaptive parameter variation scheme for metaheuristics based on the epsilon-constraint method, European Journal of Operational Research 169 (2006) 932-942. doi:10.1016/j.ejor.2004.08.029.
- [25] K. Deb, H. Jain, An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints, IEEE Transactions on Evolutionary Computation 18 (2013) 577–601. doi:10.1109/TEVC.2013.2281535.
- [26] H. Jain, K. Deb, An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part II: Handling constraints and extending to an adaptive approach, IEEE Transactions on Evolutionary Computation 18 (2013) 602-622. doi:10.1109/TEVC. 2013.2281534.
- [27] C. Audet, G. Savard, W. Zghal, A mesh adaptive direct search algorithm for multiobjective optimization, European Journal of Operational Research 204 (2010) 545-556. doi:10.1016/j.ejor.2009.11.010.
- [28] J. Knowles, ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems, IEEE Transactions on Evolutionary Computation 8 (2006) 1341–66. doi:10.1109/tevc.2005.851274.
- [29] G. Cocchi, G. Liuzzi, A. Papini, M. Sciandrone, An implicit filtering algorithm for derivative-free multiobjective optimization with box constraints, Computational Optimization and Applications 69 (2018) 267–296. doi:10.1007/s10589-017-9953-2.
- [30] G. Liuzzi, S. Lucidi, F. Rinaldi, A derivative-free approach to constrained multiobjective nonsmooth optimization, SIAM Journal on Optimization 26 (2016) 2744–2774. doi:10.1137/15M1037810.
- [31] C. Audet, J. Bigeon, D. Cartier, S. Le Digabel, L. Salomon, Performance indicators in multiobjective optimization, European Journal of Operational Research 292 (2021) 397-422. doi:10.1016/j.ejor. 2020.11.016.
- [32] T. H. Chang, L. T. Watson, J. Larson, N. Neveu, W. I. Thacker, S. Deshpande, T. C. H. Lux, Algorithm 1028: VTMOP: Solver for blackbox multiobjective optimization problems, ACM Transactions on Mathematical Software 48 (2022) Article No. 36. doi:10.1145/3529258.

- [33] K. Bringmann, T. Friedrich, Approximation quality of the hypervolume indicator, Artificial Intelligence 195 (2013) 265–290. doi:10.1016/j.artint.2012.09.005.
- [34] K. Shang, H. Ishibuchi, L. He, L. M. Pang, A survey on the hypervolume indicator in evolutionary multiobjective optimization, IEEE Transactions on Evolutionary Computation 25 (2020) 1–20. doi:10. 1109/TEVC.2020.3013290.
- [35] S. Bandyopadhyay, S. Saha, U. Maulik, K. Deb, A simulated annealing-based multiobjective optimization algorithm: AMOSA, IEEE Transactions on Evolutionary Computation 12 (2008) 269–283. doi:10.1109/TEVC.2007.900837.
- [36] J. Larson, M. Menickelly, S. M. Wild, Derivative-free optimization methods, Acta Numerica 28 (2019) 287–404. doi:10.1017/S0962492919000060.
- [37] E. Bradford, A. M. Schweidtmann, A. Lapkin, Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm, Journal of Global Optimization 71 (2018) 407–438. doi:10.1007/s10898-018-0609-2.
- [38] J. Müller, SOCEMO: Surrogate optimization of computationally expensive multiobjective problems, INFORMS Journal on Computing 29 (2017) 581–596. doi:10.1287/ijoc.2017.0749.
- [39] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, J. R. Martins, A Python surrogate modeling framework with derivatives, Advances in Engineering Software 135 (2019) 102-662. doi:10. 1016/j.advengsoft.2019.03.005.
- [40] E. W. Cheney, W. A. Light, A Course in Approximation Theory, Graduate Studies in Mathematics, AMS, Providence, RI, USA, 2009.
- [41] A. R. Conn, K. Scheinberg, L. N. Vicente, Geometry of interpolation sets in derivative free optimization, Mathematical programming 111 (2008) 141–172. doi:10.1007/s10107-006-0073-5.
- [42] A. R. Conn, K. Scheinberg, L. N. Vicente, Introduction to derivative-free optimization, MPS-SIAM Series on Optimization, SIAM, Philadelphia, PA, USA, 2009. doi:10.1137/1.9780898718768.
- [43] J. Thomann, G. Eichfelder, A trust-region algorithm for heterogeneous multiobjective optimization, SIAM Journal on Optimization 29 (2019) 1017–1047. doi:10.1137/18m1173277.
- [44] M. Berkemeier, S. Peitz, Derivative-free multiobjective trust region descent method using radial basis function surrogate models, Mathematical and Computational Applications 26 (2021) 31. doi:10.3390/ mca26020031.
- [45] J.-H. Ryu, S. Kim, A derivative-free trust-region method for biobjective optimization, SIAM Journal on Optimization 24 (2014) 334–362. doi:10.1137/120864738.
- [46] T. M. Ragonneau, Z. Zhang, PDFO: Cross-platform interfaces for Powell's derivative-free optimization solvers, 2021. URL: https://github.com/pdfo/pdfo.
- [47] J. Larson, S. M. Wild, Asynchronously parallel optimization solver for finding multiple minima, Mathematical Programming Computation 10 (2018) 303–332. doi:10.1007/s12532-017-0131-4.
- [48] R. Garnett, Bayesian Optimization, Cambridge University Press, 2023. URL: https://bayesoptbook. com.
- [49] M. Emmerich, K. Yang, A. Deutz, H. Wang, C. M. Fonseca, A multicriteria generalization of Bayesian global optimization, in: P. M. Pardalos, A. Zhigljavsky, J. Žilinskas (Eds.), Advances in Stochastic and Deterministic Global Optimization, Springer, 2016, pp. 229–242. doi:10.1007/978-3-319-29975-4.
- [50] P. Feliot, J. Bect, E. Vazquez, A Bayesian approach to constrained single- and multi-objective optimization, Journal of Global Optimization 67 (2016) 97–133. doi:10.1007/s10898-016-0427-3.

- [51] S. Daulton, M. Balandat, E. Bakshy, Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, volume 33, Curran Associates, Inc., 2020, pp. 9851-9864. URL: https://proceedings.neurips.cc/paper/2020/file/ 6fec24eac8f18ed793f5eaad3dd7977c-Paper.pdf.
- [52] D. Eriksson, M. Pearce, J. Gardner, R. D. Turner, M. Poloczek, Scalable global optimization via local bayesian optimization, Advances in Neural Information Processing Systems 32 (2019) 1-12. URL: https://proceedings.neurips.cc/paper/2019/file/ 6c990b7aca7bc7058f5e98ea909e924b-Paper.pdf.
- [53] S. S. Garud, I. A. Karimi, M. Kraft, Smart sampling algorithm for surrogate model development, Computers & Chemical Engineering 96 (2017) 103–114. doi:10.1016/j.compchemeng.2016.10.006.
- [54] P. T. Roy, A. B. Owen, M. Balandat, M. Haberland, Quasi-monte carlo methods in python, Journal of Open Source Software 8 (2023) 5309. doi:10.21105/joss.05309.
- [55] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A. G. Wilson, E. Bakshy, BoTorch: A framework for efficient Monte-Carlo Bayesian optimization, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, volume 33, Curran Associates, Inc., 2020, pp. 21524-21538. URL: https://proceedings.neurips.cc/paper/2020/file/ f5b1b89d98b7286673128a5fb112cb9a-Paper.pdf.
- [56] E. F. Campana, M. Diez, G. Liuzzi, S. Lucidi, R. Pellegrini, V. Piccialli, F. Rinaldi, A. Serani, A multiobjective DIRECT algorithm for ship hull optimization, Computational Optimization and Applications 71 (2018) 53–72. doi:10.1007/s10589-017-9955-0.
- [57] J. Bigeon, S. Le Digabel, L. Salomon, DMulti-MADS: Mesh adaptive direct multisearch for blackbox multiobjective optimization, Computational Optimization and Applications 79 (2020) 301–338. doi:10. 1007/s10589-021-00272-9.
- [58] A. L. Custódio, J. F. A. Madeira, A. I. F. Vaz, L. N. Vicente, Direct multisearch for multiobjective optimization, SIAM Journal on Optimization 21 (2011) 1109–1140. doi:10.1137/10079731x.
- [59] A. L. Custódio, J. F. A. Madeira, MultiGLODS: global and local multiobjective optimization using direct search, Journal of Global Optimization 72 (2018) 323–345. doi:10.1007/s10898-018-0618-1.
- [60] A. Abraham, L. Jain, R. Goldberg (Eds.), Evolutionary Multiobjective Optimization: Theoretical Advances and Applications, Advanced Information and Knowledge Processing Series, Springer Verlag, London, UK, 2005. doi:10.1007/1-84628-137-7.
- [61] K. Deb, A. Pratap, S. Agarwel, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Transactions on Evolutionary Computation 6 (2002) 182–197. doi:10.1109/4235. 996017.
- [62] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nature Reviews Physics 3 (2021) 422–440. doi:10.1038/s42254-021-00314-5.
- [63] E. Pickering, S. Guth, G. E. Karniadakis, T. P. Sapsis, Discovering and forecasting extreme events via active learning in neural operators, Nature Computational Science 2 (2022) 823–833.
- [64] R. Astudillo, P. I. Frazier, Thinking inside the box: a tutorial on grey-box bayesian optimization, in: Proc. 2021 Winter Simulation Conference (WSC 2021), IEEE, 2021. doi:10.1109/WSC52266.2021. 9715343.
- [65] K. A. Khan, J. Larson, S. M. Wild, Manifold sampling for optimization of nonconvex functions that are piecewise linear compositions of smooth components, SIAM Journal on Optimization 28 (2018) 3001–3024. doi:10.1137/17m114741x.

- [66] J. Larson, M. Menickelly, Structure-aware methods for expensive derivative-free nonsmooth composite optimization, Mathematical Programming Computation 16 (2024) 1–36. doi:10.1007/ s12532-023-00245-5.
- [67] H. Zhang, A. R. Conn, On the local convergence of a derivative-free algorithm for least-squares minimization, Computational Optimization and Applications 51 (2012) 481–507. doi:10.1007/ s10589-010-9367-x.
- [68] S. M. Wild, Solving derivative-free nonlinear least squares problems with POUNDERS, in: T. Terlaky, M. F. Anjos, S. Ahmed (Eds.), Advances and Trends in Optimization with Engineering Applications, SIAM, 2017, pp. 529–540. doi:10.1137/1.9781611974683.ch40.
- [69] A. Benítez-Hidalgo, A. J. Nebro, J. García-Nieto, I. Oregi, J. Del Ser, jMetalPy: A Python framework for multi-objective optimization with metaheuristics, Swarm and Evolutionary Computation 51 (2019) 100598. doi:10.1016/j.swevo.2019.100598.
- [70] F. Biscani, D. Izzo, A parallel global multiobjective framework for optimization: pagmo, Journal of Open Source Software 5 (2020) 2338. doi:10.21105/joss.02338.
- [71] J. Blank, K. Deb, pymoo: Multi-objective optimization in Python, IEEE Access 8 (2020) 89497–89509. doi:10.1109/ACCESS.2020.2990567.
- [72] S. Le Digabel, Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm, ACM Transactions on Mathematical Software 37 (2011) Article No. 44. doi:10.1145/1916461.1916468.
- [73] G. Feldman, S. R. Hunter, SCORE allocations for bi-objective ranking and selection, ACM Transactions on Modeling Computer and Simulation 28 (2018). doi:10.1145/3158666.
- [74] C. Audet, E. Hallé-Hannan, S. Le Digabel, A general mathematical framework for constrained mixedvariable blackbox optimization problems with meta and categorical variables, in: Operations Research Forum, volume 4, Springer, 2023, p. 12. doi:10.1007/s43069-022-00180-6.
- [75] P. Saves, R. Lafage, N. Bartoli, Y. Diouane, J. Bussemaker, T. Lefebvre, J. T. Hwang, J. Morlier, J. R. R. A. Martins, SMT 2.0: A Surrogate Modeling Toolbox with a focus on Hierarchical and Mixed Variables Gaussian Processes, Technical Report, arXiv preprint, 2023. doi:10.48550/arXiv.2305. 13998.
- [76] H. Moriwaki, Y.-S. Tia, N. Kawashita, T. Takagi, Mordred: a molecular descriptor calculator, Journal of Cheminformatics 10 (2018). doi:10.1186/s13321-018-0258-y.
- [77] S. Le Digabel, S. M. Wild, A Taxonomy of Constraints in Black-Box Simulation-Based Optimization, Technical Report 2, 2024. doi:10.1007/s11081-023-09839-3.
- [78] C. Audet, J. E. Dennis, A progressive barrier for derivative-free nonlinear programming, SIAM Journal on Optimization 20 (2009) 445–472. doi:10.1137/070692662.
- [79] R. M. Lewis, V. Torczon, A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds, SIAM Journal on Optimization 12 (2002) 1075–1089. doi:10.1137/S1052623498339727.
- [80] G. Cocchi, M. Lapucci, An augmented Lagrangian algorithm for multi-objective optimization, Computational Optimization and Applications 77 (2020) 29–56. doi:10.1007/s10589-020-00204-z.
- [81] R. Chard, Y. Babuji, Z. Li, T. Skluzacek, A. Woodard, B. Blaiszik, I. Foster, K. Chard, funcX: A federated function serving fabric for science, in: Proc. 29th International Symposium on High-Performance Parallel and Distributed Computing (HPDC '20), ACM, 2020, pp. 65–76. doi:10.1145/ 3369583.3392683.
- [82] S. Hudson, J. Larson, J.-L. Navarro, S. M. Wild, libEnsemble: A complete Python toolkit for dynamic ensembles of calculations, Technical Report 92, 2023. doi:10.21105/joss.06031.

- [83] S. Hudson, J. Larson, J.-L. Navarro, S. M. Wild, libEnsemble: A library to coordinate the concurrent evaluation of dynamic ensembles of calculations, IEEE Transactions on Parallel and Distributed Systems 33 (2022) 977–988. doi:10.1109/tpds.2021.3082815.
- [84] R. M. Kolonay, M. Sobolewski, Service oriented computing environment (SORCER) for large scale, distributed, dynamic fidelity aeroelastic analysis, in: International Forum on Aeroelasticity and Structural Dynamics (IFASD 2011), Optimization, Citeseer, 2011, pp. 26–30. URL: http: //citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.656.7539.
- [85] T. H. Chang, J. Larson, L. T. Watson, T. C. H. Lux, Managing computationally expensive blackbox multiobjective optimization problems using libEnsemble, in: Proc. 2020 Spring Simulation Conference (SpringSim 2020), the 28th High Performance Computing Symposium (HPC '20), SCS, 2020, p. Article No. 31. doi:10.22360/SpringSim.2020.HPC.001.
- [86] C. Raghunath, T. H. Chang, L. T. Watson, M. Jrad, R. K. Kapania, R. M. Kolonay, Global deterministic and stochastic optimization in a service oriented architecture, in: Proc. 2017 Spring Simulation Conference (SpringSim 2017), the 25th High Performance Computing Symposium (HPC '17), SCS, Virginia Beach, VA, USA, 2017, p. Article No. 7. doi:10.22360/springsim.2017.hpc.023.
- [87] M. A. Heroux, L. McInnes, D. E. Bernholdt, A. Dubey, E. Gonsiorowski, O. Marques, J. D. Moulton, B. Norris, E. Raybourn, S. Balay, R. A. Bartlett, L. Childers, T. Gamblin, P. Grubel, R. Gupta, R. Hartman-Baker, J. C. Hill, S. Hudson, C. Junghans, A. Klinvex, R. Milewicz, M. Miller, H. Ah Nam, J. O'Neal, K. Riley, B. Sims, J. Schuler, B. F. Smith, L. Vernon, G. R. Watson, J. Willenbring, P. Wolfenbarger, Advancing Scientific Productivity through Better Scientific Software: Developer Productivity and Software Sustainability Report, Technical Report ORNL TM-2020 1459 / ECP-U-RPT-2020-0001, Oak Ridge National Laboratory, Oak Ridge, TN, USA, 2020. doi:10.2172/1606662.
- [88] C. Audet, S. Le Digabel, V. Rochon Montplaisir, C. Tribes, Algorithm 1027: NOMAD version 4: Nonlinear optimization with the MADS algorithm, ACM Transactions on Mathematical Software 48 (2022). doi:10.1145/3544489.
- [89] S. Mannor, V. Perchet, G. Stoltz, Approachability in unknown games: Online learning meets multi-objective optimization, in: Proc. 27th Conference on Learning Theory (PMLR), volume 35 of *Proceedings of Machine Learning Research*, PMLR, Barcelona, Spain, 2014, pp. 339–355. URL: https://proceedings.mlr.press/v35/mannor14.html.
- [90] C. F. Hayes, R. Rădulescu, E. Bargiacchi, J. Källström, M. Macfarlane, M. Reymond, T. Verstraeten, L. M. Zintgraf, R. Dazeley, F. Heintz, et al., A practical guide to multi-objective reinforcement learning and planning, Autonomous Agents and Multi-Agent Systems 36 (2022) 1–59. doi:10.1007/ s10458-022-09552-y.
- [91] T. P. Sapsis, A. Blanchard, Optimal criteria and their asymptotic form for data selection in datadriven reduced-order modelling with Gaussian process regression, Philosophical Transactions of the Royal Society A 380 (2022) 20210197. doi:10.1098/rsta.2021.0197.
- [92] E. Zitzler, M. Laumanns, L. Thiele, SPEA2: Improving the strength Pareto evolutionary algorithm, TIK-report 103 (2001). doi:10.3929/ethz-a-004284029.
- [93] D. Hadka, Platypus multiobjective optimization in Python, Technical Report Version 1.0.4, GitHub, 2015. URL: https://platypus.readthedocs.io/en/latest.
- [94] Y. Tian, R. Cheng, X. Zhang, Y. Jin, PlatEMO: A MATLAB platform for evolutionary multi-objective optimization [educational forum], IEEE Computational Intelligence Magazine 12 (2017) 73–87. doi:10. 1109/MCI.2017.2742868.
- [95] J. J. Durillo, A. J. Nebro, jMetal: A Java framework for multi-objective optimization, Advances in Engineering Software 42 (2011) 760-771. doi:10.1016/j.advengsoft.2011.05.014.

- [96] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, C. Gagné ", DEAP: Evolutionary algorithms made easy, Journal of Machine Learning Research 13 (2012) 2171-2175. URL: https: //www.jmlr.org/papers/v13/fortin12a.html.
- [97] G. e. a. Liuzzi, Dfo-lib, 2024. URL: https://github.com/DerivativeFreeLibrary.
- [98] S. Tavares, C. P. Brás, A. L. Custódio, V. Duarte, P. Medeiros, Parallel strategies for direct multisearch, Numerical Algorithms 92 (2022) 1757–1788. doi:10.1007/s11075-022-01364-1.
- [99] K. Cooper, S. R. Hunter, PyMOSO: Software for multi-objective simulation optimization with R-PERLE and R-MinRLE, INFORMS Journal on Computing 32 (2020) 1101–1108. doi:10.1287/ijoc. 2019.0902.
- [100] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, highperformance deep learning library, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 32, Curran Associates, Inc., 2019, pp. 1–12. URL: https://proceedings.neurips.cc/paper/2019/file/ bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- [101] T. H. Chang, S. M. Wild, H. Dickinson, ParMOO: Python library for parallel multiobjective simulation optimization, Technical Report Version 0.3.1, Argonne National Laboratory, Lemont, IL, USA, 2023. URL: https://parmoo.readthedocs.io/en/latest.
- [102] E. Gamma, R. Helm, R. Johnson, J. Vlissides, Design patterns: elements of reusable object-oriented software, Addison-Wesley, Reading, MA, USA, 1995.
- [103] I. Dunning, J. Huchette, M. Lubin, JuMP: A modeling language for mathematical optimization, SIAM Review 59 (2017) 295–320. doi:10.1137/15M1020575.
- [104] W. E. Hart, C. D. Laird, J.-P. Watson, D. L. Woodruff, G. A. Hackebeil, B. L. Nicholson, J. D. Siirola, Pyomo – optimization modeling in Python, Springer Optimization and Its Applications, 2 ed., Springer Cham, Cham, Switzerland, 2017. doi:10.1007/978-3-319-58821-6.
- [105] G. Chen, T. H. Chang, J. Power, C. Jing, An integrated multi-physics optimization framework for particle accelerator design, in: Proc. 2023 Winter Simulation Conference (WSC 2023), Industrial Applications Track, 2023. doi:10.48550/arXiv.2311.09415.
- [106] K. Deb, L. Thiele, M. Laumanns, E. Zitzler, Scalable multi-objective optimization test problems, in: Proc. 2002 IEEE Congress on Evolutionary Computation (CEC '02), volume 1, IEEE, 2002, pp. 825–830. doi:10.1109/CEC.2002.1007032.
- [107] F. Chollet, et al., Keras, https://keras.io, 2015.
- [108] C. Zhu, R. H. Byrd, P. Lu, J. Nocedal, Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization, ACM Transactions on Mathematical Software 23 (1997) 550–560. doi:10.1145/279232.279236.
- [109] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, S. Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nature Methods 17 (2020) 261–272. doi:10.1038/s41592-019-0686-2.

- [110] T. H. Chang, J. R. Elias, S. M. Wild, S. Chaudhuri, J. A. Libera, A framework for fully autonomous design of materials via multiobjective optimization and active learning: challenges and next steps, in: 11th Intl. Conf. on Learning Representation (ICLR 2023), Workshop on Machine Learning for Materials (ML4Materials), 2023, pp. 1–10. URL: https://openreview.net/forum?id=8KJS7RPjMqG, to appear.
- [111] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825-2830. URL: https://www.jmlr.org/papers/volume12/ pedregosa11a/pedregosa11a.pdf.
- [112] T. H. Chang, S. M. Wild, Designing a framework for solving multiobjective simulation optimization problems, 2024. doi:10.1287/ijoc.2023.0250.cd, available for download at https://github.com/INFORMSJoC/2023.0250.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. http://energy.gov/downloads/doe-public-access-plan