Report from Dagstuhl Seminar 23361

Multiobjective Optimization on a Budget

Richard Allmendinger*[1], Carlos M. Fonseca*[2], Serpil Sayin*[3], Margaret M. Wiecek*[4], and Michael Stiglmayr†[5]

1 **University of Manchester, GB.** richard.allmendinger@manchester.ac.uk
2 **University of Coimbra, PT.** cmfonsec@dei.uc.pt
3 **Koc University – Istanbul, TR.** ssayin@ku.edu.tr
4 **Clemson University, US.** wmalgor@clemson.edu
5 **Universität Wuppertal, DE.** stiglmayr@uni-wuppertal.de

Abstract

The Dagstuhl Seminar 23361 Multiobjective Optimization on a Budget carried on a series of seven previous Dagstuhl Seminars (04461, 06501, 09041, 12041, 15031, 18031, 20031) focused on Multiobjective Optimization. The original goal of this series has been to strengthen the links between the Evolutionary Multiobjective Optimization (EMO) and the Multiple Criteria Decision Making (MCDM) communities, two of the largest communities concerned with multiobjective optimization today. This seminar particularly focused on the case where the approaches from both communities may be challenged by limited resources.

This report documents the program and the outcomes of Dagstuhl Seminar 23361 "Multiobjective Optimization on a Budget". Three major types of resource limitations were highlighted during the seminar: methodological, technical and human related. The effect of these limitations on optimization and decision-making quality, as well as methods to quantify and mitigate this influence, were considered in different working groups.

Seminar September 3–8, 2023 – http://www.dagstuhl.de/23361

2012 ACM Subject Classification Computing methodologies → Machine learning; Computing methodologies → Neural networks; Mathematics of computing → Evolutionary algorithms; Theory of computation → Mathematical optimization; Computing methodologies → Computational control theory

Keywords and phrases evolutionary algorithms, expensive optimization, few-shot learning, machine learning, optimization, simulation

Digital Object Identifier 10.4230/DagRep.13.9.1

1 Executive Summary

Richard Allmendinger (University of Manchester, GB)
Carlos M. Fonseca (University of Coimbra, PT)
Serpil Sayin (Koc University – Istanbul, TR)
Margaret M. Wiecek (Clemson University, US)

Multiobjective optimization (MO), a discipline within systems science that provides models, theories, and methodologies to address decision-making problems under conflicting objectives, has a myriad of applications in all areas of human activity ranging from business and management to engineering. This seminar is a result of the desire to continue to make

* Editor / Organizer
† Editorial Assistant / Collector

Multiobjective Optimization on a Budget, *Dagstuhl Reports*, Vol. 13, Issue 9, pp. 1–68
Editors: Richard Allmendinger, Carlos M. Fonseca, Serpil Sayin, and Margaret M. Wiecek
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

MO useful to society as it faces complex decision-making problems and experiences limited resources for decision making. Of particular interest are processes that evolve competitively in environments with scarce resources and lead to decision problems that are characterized by multiple, incommensurate, and conflicting objectives, and engage multiple decision-makers. Viewing optimization and decision making as the complementary aspects of the multiobjective paradigm, the seminar set out to focus around three major types of resource limitations: methodological (e.g., number of solution evaluations), technical (e.g., computation time, energy consumption), and human related (e.g., decision maker availability and responsiveness). The effect of these limitations on optimization and decision-making quality, as well as methods to quantify and mitigate this influence, were of particular interest. Ideas related to modelling, theory, algorithm design, benchmarking, performance metrics, and novel applications of MO under budget constraints were discussed.

To initiate a discussion among the participants on how to address challenges of MO under a budget, the organizers presented specific research directions at the beginning of the seminar. These directions along with their highlights are described below.

- Model reduction: In the MO problem not all functions may be of interest to the decision maker (DM) or not all objectives may be in conflict with each other. Under a limited budget, it is of interest to make the original problem simpler by removing unnecessary objective functions while the solution set remains unchanged. Another reason to reduce the problem is its size. MO problems with four or more criteria bring computational and decision-making challenges that are not typical when the number of objectives is lower.
- Model decomposition and coordination-based decision making: If a reduction of the objectives is not possible, then the solution of the overall MO problem in its entirety may be challenging or even impossible to obtain. In this situation, decomposition of the MO problem into a set of MO subproblems with a smaller number of criteria becomes appealing provided solving the subproblems can be coordinated and related to solving the original problem. When the MO problem is decomposed while computation of the overall solution set is possible, the decomposition goal is to enhance capability of making coordinated tradeoff decisions by working in lower dimensional spaces, which decreases the cognitive burden on DMs. Otherwise, if computation of the overall solution set is not possible, the decomposition goal becomes more challenging since the intention is to coordinate the subproblems' solution sets to construct the overall set and to facilitate decision making in a similar way.
- Representation of the optimization solution set: It is of interest to design cost-effective methods for obtaining a complete or partial description of the Pareto set. An exact description of this set might be available analytically as a closed-form formula, numerically as a set of points, or in mixed form as a parametrized set of points. Unfortunately, for the majority of MO problems, it is not easy to obtain an exact description of the solution set that includes typically a very large number or infinite number of points. Even if it is theoretically possible to find these points exactly, this is often computationally challenging and expensive, and therefore is usually abandoned. On the other hand, if it is possible to obtain the complete solution set, one might not be interested in this task due to overflow of information. Another reason for approximating the solution set, rather than finding the solution set exactly, is that many real-world problems (e.g., in engineering) cannot be completely and correctly formulated before a solution procedure starts. Since the exact solution set is very often not attainable, an approximated description of the solution set becomes an appealing alternative.

- Surrogate-assisted optimization: The combination of evolutionary MO (EMO) algorithms with efficient computational models, often known as metamodels or surrogates, has become a common approach to approximate outcomes of a time-consuming, expensive, and/or resource intense simulation or physical experiment, and thus to tackle problems with a limited budget. Surrogate-assisted (SA) methods vary in aspects such as the use of the metamodel (e.g., different models for different objective functions or one model for all objective functions), type of metamodel (e.g., Gaussian process, radial basis neural network, etc.), how the metamodel is updated (e.g., expected improvement, expected hypervolume improvement), and training time of the metamodel. In particular, the combination of optimization with Gaussian process approximation, known as Bayesian optimization, is a recent trend to efficiently deploy data in model development.
- Multistage optimization: In real-world applications, problem data does not always become available all at once, but at different points in time until a final decision needs to be made. In particular, waiting until all the required data is available may not leave enough time to run the optimization process on the whole problem and successfully compute a final decision. In addition, it is often possible to model the uncertainty associated with the yet unknown data given the data that is already known, at least to some extent. Two-stage (and, more generally, multi-stage) approaches to optimization reformulate the original problem as a number of sub-problems to be solved sequentially, in such a way that the last problem(s) in the sequence can effectively be solved in the (short) time available.
- Preference acquisition and communication with the decision maker: The ultimate goal in MO is to serve one or multiple DMs whose goal is to come up with a single most preferred solution from among the ones that are available. Given an optimization model, DM's preferences may be incorporated prior to, during or after employing a solution procedure. In particular, interactive methods require the DM's involvement in the solution process during which they reveal their preferences based on the presented information. Under a limited budget, communication with the DM shall be designed effectively and economically.
- Benchmarking of algorithms: SA methods are considered as the method of choice to tackle problems subject to a limited budget in terms of function evaluations. However, SA methods are not often compared to widely different alternatives (e.g., different kernels and distance measures, non-SA methods, etc.), and are often tested on narrow sets of problems (multimodal, low-dimensional, static, deterministic, unconstrained, and continuous functions) and rarely on real-world problems, which makes it difficult to assess where (or if) these methods actually achieve state-of-the-art performance in practice. Moreover, several aspects in the design of SA algorithms vary across implementations without a clear recommendation emerging from current practices, and many of these design choices are not backed up by authoritative test campaigns. This seminar topic aimed to raise awareness and hence a push to more work being carried out on developing benchmarking guidelines for SA algorithms.

In response to the presented research directions, some participants found research topics of interest among those suggested by the organizers. These topics included model reduction, decomposition and coordination, solution set representation, and surrogate modeling. Other participants proposed different topics that also targeted the theme of MO under a budget. Those topics included design of experiments for MO, correlations in MO, and design of evolutionary algorithms. Overall, seven research topics were proposed and pursued.

Independently of developing and forming research topics, a collection of eight talks were given during the seminar. Two of the speakers were considered "invited" because they were asked before the seminar to give a talk. These talks addressed two of the research directions initiated by the organizers. The other speakers, being inspired by the ongoing seminar, proposed talks that were integrated daily into the seminar program. The invited and contributed talks kept the seminar in balance ensuring ample time for working in groups.

During the seminar the schedule was updated on a daily basis to maintain flexibility in balancing time slots for the invited and contributed talks, discussions, and working group sessions. The working groups were established on the first day in an interactive fashion. Starting with three large working groups focused around the three central topics of the seminar (methodological, technical, and human-related resource limitations), participants were invited to formulate their favorite topics and most important challenges. The three initial groups split to eventually form eight groups by the end of the seminar. During the week the participants were allowed to change the working groups based on their research interest. The abstracts of the delivered talks and the extended abstracts of the working groups can be found in the subsequent chapters of this report.

Further notable events during the week included: (i) a hike that took place on Wednesday afternoon, (ii) a session allowing the participants to share the details of upcoming professional events in the research community, (iii) a joint session with the participants of the concurrent seminar 23362 "Decision-Making Techniques for Smart Semiconductor Manufacturing" and (iv) an informal get together on Thursday evening.

Offers and Needs Market

An *Offers* & *Needs Market* ran throughout the entire week. The participants could write their research offers and needs regarding MO on note paper in different colors and post them on pin boards (see Fig. 1) to attract or find a possible collaborator. Participants discussed potential collaboration opportunities during the coffee breaks and after hours.

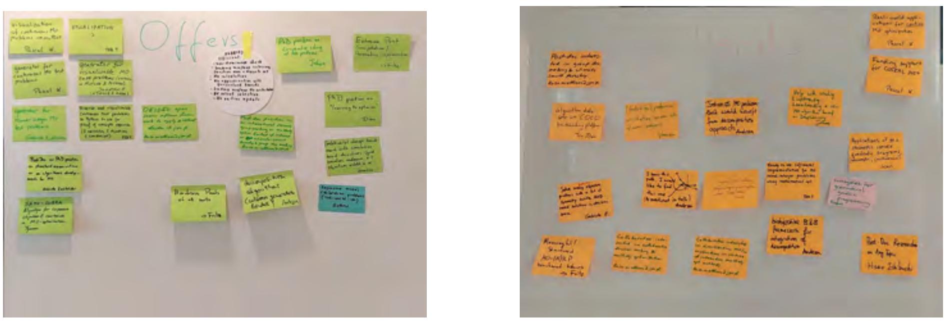

Figure 1 Offers and needs market.

Outcomes

The outcomes of each of the working groups can be seen in the sequel.

The organizers have arranged a special issue of the *Journal of Multi-Criteria Decision Analysis* entitled "Multiobjective Optimization on a Budget" for which they will serve as Guest Editors. This issue will be an outlet for papers authored and submitted by the seminar's participants as well as by researchers world-wide.

This seminar resulted in a very insightful, productive and enjoyable week. It has already led to first new results, cooperations and research topics.

Acknowledgements

The organizers would like to express their appreciation to the Dagstuhl Office and its helpful and patient staff for their professional support and smooth cooperation; huge thanks to the organizers of the previous seminars in this series for setting us up for success; and thanks to all the participants, who worked hard and were amiable company all week.

In a later section, we also give special thanks to Margaret Wiecek as she steps down from the organizer role.

2 Table of Contents

3.9 Multi-objective Branch-and-Bound on a Budget

Michael Stiglmayr (Universität Wuppertal, DE)

Joint work of Julius Bauß, Michael Stiglmayr

In this talk we discuss modifications of multi-objective branch-and-bound to diversify solutions and yield a good approximation of the non-dominated set when only limited computation time is available. It is crucial not only to find efficient solutions in early stages of the algorithm but also to find a set of solutions whose images are close to and well distributed along the non-dominated frontier. In particular the adaptation of branching and queuing of sub-problems seems to be important. We use, e.g., the hypervolume indicator as a measure for the gap between lower and upper bound set to implement a multi-objective best-first strategy. Moreover, gap measure indicate the solution quality when prematurely stopping the branch-and-bound algorithm.

References

1 Bauß, J., Stiglmayr, M.: Augmenting bi-objective branch and bound by scalarization-based information (2023), `https://arxiv.org/abs/2301.11974`
2 Bauß, J., Stiglmayr, M.: Adapting branching and queuing for multi-objective branch and bound (2023), `https://arxiv.org/abs/2311.05980`

4 Working groups

4.1 Decoupled Design of Experiments for Multi-objective Optimisation on a Budget

Mickaël Binois (INRIA – Sophia Antipolis, FR), Jürgen Branke (University of Warwick, GB), Jonathan Fieldsend (University of Exeter, GB), Robin Purshouse (University of Sheffield, GB)

4.1.1 Introduction

Fundamental to the performance of surrogate-based optimisation frameworks is the need to construct an *initial model* based on a carefully selected set of initial designs, and any prior system knowledge. This is both in the case of Bayesian optimisation, which used and iteratively update model(s) mapping decision vectors to predicted performance criteria values, and for evolutionary computation approaches which involving surrogates. The selection and construction of initial designs, which are often treated separately to the decision vectors queried during the subsequent optimisation process, are usually referred to as the *design of experiments* (or DoE for short). This is because these decision vectors are selected to – in some fashion – be maximally informative on the global underlying process, rather than being biased towards particular regions.

Without any prior information regarding the properties of the objective function(s) such DoE for model fitting are commonly based around *space filling* sequences such as Latin hypercube sampling [9] or Sobol sequences [10], as purely random sampling tends to naturally result in clusters, which do not serve model fitting well, particularly when the budget for sampling is tight.

Where there are multiple criteria being modelled, this leads to an interesting and under-explored question: *should one evaluate all initial designs fully, or selectively evaluate a subset of objectives per design, allowing a greater number of locations to be partially evaluated when building the model(s)?* A few works have looked at *decoupling* objective evaluations during the search process – particularly where there are different costs associated with each objective, but this can also be advantageous where there is a difference in the complexity of the functions being modelled (e.g. one being smooth slowly changing, the other being rugged and fast changing). As such, this appears to be a promising direction for further investigation and research, as even small improvements in such areas can effectively lead to large savings for expensive optimisation problems.

4.1.2 Related Work

A small number of existing works have considered decoupled and/or cost-aware multi-objective optimisation – some of which have considered these factors during the initial DoE phase. Below we discuss the most relevant approaches. A wider survey on the topic of objectives with different costs can be found in [1].

Hernández-Lobato and colleagues proposed the *Predictive Entropy Search for Multi-Objective Bayesian Optimization* (PESMO) method [6]. PESMO uses predictive entropy search as the acquisition function. This function represents each objective using an additive component, which enables a decoupled evaluation approach to be adopted. The approach was subsequently extended to also consider constraints (again where decoupling is possible) [5].

Suzuki et al. developed the *Pareto-frontier entropy search* (PFES) approach [11]. PFES is also an entropy approach but considers the entropy in objective-space rather than decision-space, which is computationally simpler. This method also includes cost in evaluating the objectives by including cost in the denominator of the acquisition function. Like PESMO, the approach is easily extended to consider decoupled evaluations.

Iqbal and colleagues proposed the *Flexible Multi-Objective Bayesian Optimization* (FlexiBo) algorithm [7]. The approach uses a decoupled evaluation in the Bayesian optimisation run but uses a coupled initial DoE procedure. FlexiBo includes two main features: (1) a new acquisition function that is the expected change in hypervolume if only one objective function is evaluated, divided by the cost of this function evaluation; and (2) a confidence region in the objective space for the partially evaluated points. The estimated cost of evaluating each objective is updated each time the objective is evaluated – this is a mean estimate of the cost (treating any observed variability as occurring at random).

Most recently, Buckingham et al. extended the multi-attribute Knowledge Gradient [2] to the case where objectives can be evaluated independently [3]. The authors demonstrate the benefit of independent evaluation not only when the computational times for objectives differ, but also when the lengthscales of the modelled landscapes (which determine the smoothness of the landscape) differ.

A slightly different problem is considered in [8], where one objective is much cheaper (essentially free) to evaluate than the other. They directly incorporate evaluation of the cheap objectives into a pair of hypervolume-based acquisition functions for BO. Consequently, the cheap objectives are evaluated many times while the acquisition function is optimized.

A summary of the different approaches is shown in Table 1, highlighting which methods feature decoupled and cost-aware acquisition functions during the initial DoE, the subsequent optimisation run, or both phases.

Table 1 Existing methods for decoupled cost-aware multi-objective optimisation.

	Design of experiments		Optimisation		
Approach	Decoupled?	Cost-aware?	Decoupled?	Cost-aware?	Acquisition function
PESMO [6]	✓	✗	✓	✗	predictive entropy search
PFES [11]	✗	✗	✓	✓	cost-weighted Pareto frontier entropy
FlexiBO [7]	✗	✗	✓	✓	cost-weighted objective space entropy
C-MOKG [3]	✗	✗	✓	✓	cost-weighted multi-objective knowledge gradient

4.1.3 New analyses by the working group

4.1.3.1 Initial DoE when evaluations are decoupled

The costs of the objectives are assumed to be the same for now.

Goal: studying the effect on coupled vs. decoupled designs of experiments (DoE) on the uncertainty on the Pareto front.

To this end, we experiment on Gaussian process models (GPs). More precisely, we generate samples from a Gaussian process model and use it as the ground truth. The hyperparameters are supposed to be known to remove the effect of inference. Hence there is no model mismatch. Examples of outcome are given in Figure 2.

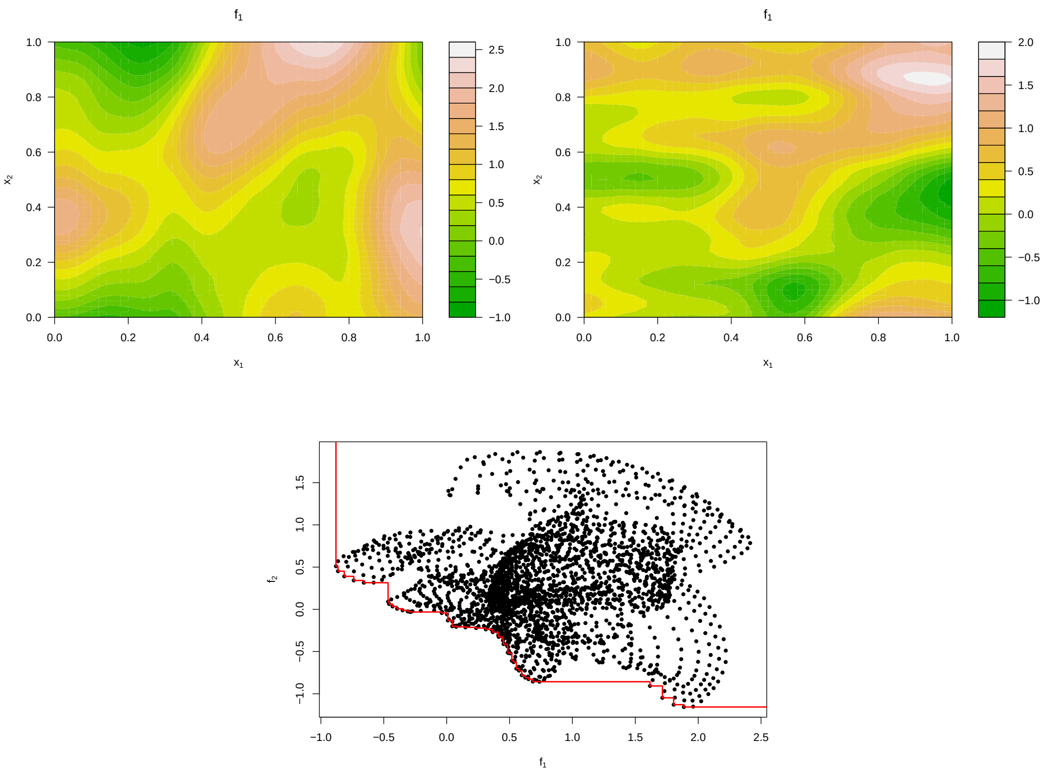

Figure 2 Top: two realisations of Gaussian process priors, with Matérn 5/2 covariance kernel, with lengthscale hyperparameters $(0.3, 0.4)$ (resp. $(0.4, 0.2)$) for f_1 (resp. f_2), and unit variance. Bottom: corresponding image in the objective space.

Next, to measure the uncertainty on the Pareto front associated with the fitted GPs, we rely on the so called Vorob'ev deviation (VD), a set based variance measure, see Algorithm 1 for a pseudo code and, e.g., [4] for the details. The reference point used for hypervolume

Algorithm 1: Pseudo-code for the testing procedure.

Generate the first design of experiments X_1 for objective 1.
(Coupled case) $X_2 = X_1$ the DoE of the second objective is the same.
(Decoupled case) Generate X_2 the second DoE.
Build GP models.
Generate s conditional samples on some designs X_s from all GPs.
Compute the s non-dominated points on couples of samples from the different GPs.
Compute the corresponding Vorob'ev deviation.

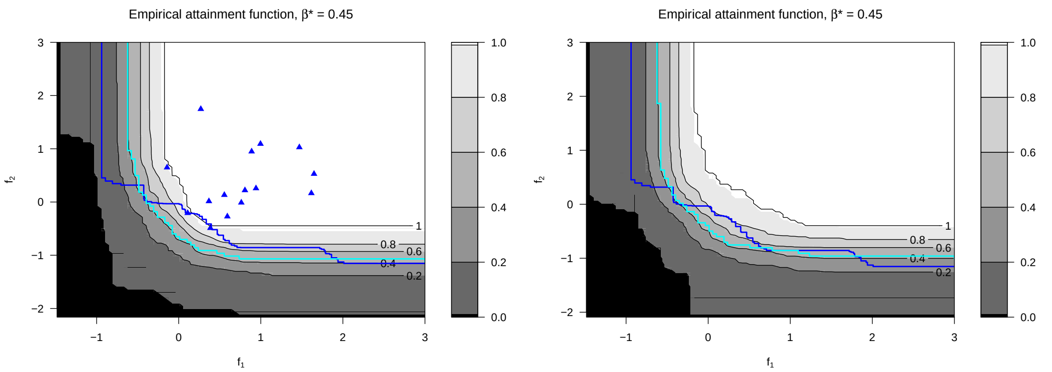

Figure 3 Attainment function representation in the coupled (left) and decoupled (right) cases. The blue triangle mark observations in the coupled case, where both objectives are evaluated. The cyan line represents the estimated Pareto front of the GP while the reference Pareto front is in blue.

computations is taken to be $(3, 3)$. An example is provided in Figure 3, where the DoE for the first objective is the same while the second one is either coupled or decoupled. One visible effect is that when both objectives are jointly evaluated, the area that is dominated (attainment value $= 1$) is larger. This is probably because in the decoupled case, solutions are never surely dominated (even though the domination probability is extremely low).

We compare VD values of different setups for the coupled and decoupled case:

- the DoE for the first objective is either uniformly sampled or is a maximin Latin hypercube design;
- the DoE for the second objective is the same as the first objective (coupled case), uniformly sampled or an LHS augmenting the DoE of the first objective.

Figure 4 shows the results. Comparing the top row (both initial designs uniformly sampled) with the middle row (an augmenting LHS used to complement the first uniform DoE), there seems to be not much difference. However, the bottom row (first design is sampled with LHS, second uses augmenting LHS) shows a significant improvement of the Vorob'ev deviation of either the coupled (red dots) or decoupled (box plots) sampling. Clearly, a space-filling design improves our estimate of the Pareto front, but it seems not sufficient to only make the design of the second objective space-filling.

Note that with respect to the Vorob'ev deviation, when at least one of the designs is random (first two rows, first two columns), the red dots are sometimes above and sometimes below the median of the boxplots, while the red dots are mostly below the median of the boxplots in the bottom row (full space-filling design). This indicates that at least if a space-filling design is used, decoupled sampling is worse than coupled sampling, possibly due to the effect mentioned above on the size of the known dominated region. Note, however, that in these experiments we assume equal cost of sampling the two objectives, and equal lengthscales of the two objectives. As we see later, in other cases decoupling may be beneficial.

The results look slightly different when considering the expected product of the standard deviations of the GP (right column), which is an indication of the accuracy of the estimation quality of the models over the entire search space, rather than the Pareto frontier. Here, the first two rows show a clear benefit of decoupled sampling. However, this benefit seems to disappear once both objectives are sampled using space-filling designs (third row).

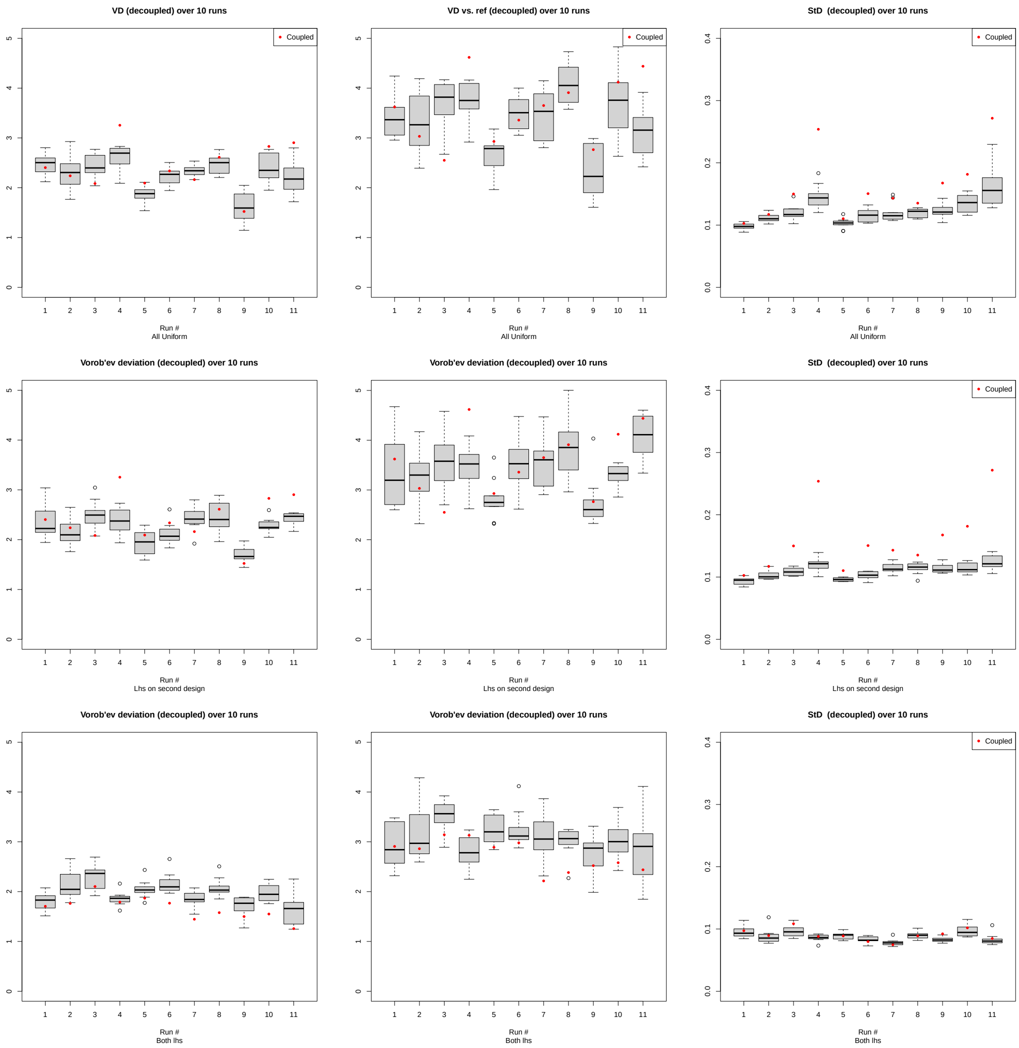

Figure 4 Boxplots of Vorob'ev deviation with decoupled designs, over 11 different runs and 10 replications per run. In the top row, both initial designs are uniformly sampled, in the middle row, an augmenting LHS is used to complement the first uniform DoE, and in the bottom row, an augmenting LHS is used to complement the first LHS design. Left column shows VD, middle column shows VD against true Pareto front, and right column shows standard deviation product. The value of the coupled design is in red.

4.1.3.2 Initial DoE when evaluations have different costs

Now let us assume the cost is different between different objectives f_1 and f_2 (etc). The first tasks are to define the total time budget for experiments and get relative costs of f_1, f_2, …, f_3. We will then consider a number of alternative approaches to DoE, including a coupled baseline.

1. (Coupled) Both functions evaluated at once.
2. (Decoupled naive) Both functions evaluated the same number of times, but at differing locations. (generated by Augmented LHS)
3. (Decoupled) The allocation of total budget to the two functions depends on lengthscales and relative costs, according to Eq. 1. Objectives with smaller lengthscales and smaller cost are sampled more often.

Considering how to split the computational budget, let us consider the simplest case of optimising a (weighted) sum of two objectives. In such a case, if we want to minimise integrated mean squared prediction error (IMSPE), then it is not possible to improve beyond coupled sampling, as the variances of the two functions just add up, and the optimal design for each function would be the same. However, if the costs or lengthscales are different, then we could use IMSPE to determine an appropriate allocation of the budget to the two functions as follows:

$$\min \frac{IMSPE(n_1)}{c_1 \times n_1} + \frac{IMSPE(N - n_1)}{c_2 \times (N - n_1)}, \tag{1}$$

where N is the total budget, n_1 is the number of samples allocated to objective f_1, and $c_1(c_2)$ are the cost of evaluating objective $f_1(f_2)$.

As in the previous section, we rely on GP samples to define a ground truth. We also assume some known values of the lengthscales of the objectives: $(0.3, 0.4)$ for the first, $(0.4, 0.2)$ for the second. We start with four initial designs for each objective in the various cases, then 26 decoupled evaluations are performed. We only compare the 'coupled', 'naive' and 'decoupled' strategies. The results are in Figure 5. First, from the IMSPE results, we observe that the values for objective 1 and 2 are different (importantly, the GP variances are equal here), due to the different lengthscales. The naive baseline always performs worst. Then, in the same cost case, there is no change between the coupled and decoupled case. As the cost of f_2 increases, the effect is that the IMSPE of f_1 is reduced faster compared to f_2, with no strong detrimental effect on f_2 for the same total cost. The outcome is that it is reasonable to sample more f_1, in a ratio that only depends on the lengthscales and relative cost.

4.1.4 Discussion and future research ideas

In this report, we have examined the possibility of improving the quality of the surrogate models obtained through a DoE in case of multi-objective optimisation where the evaluation of the different objectives can be decoupled. We found that for the case of equal lengthscales, decoupling the evaluations (i.e., evaluating different solutions on different objectives) did tend to worsen the quality of the Pareto front estimate as measured by Vorob'ev deviation. However, when objectives had different costs and/or lengthscales, decoupling could improve results substantially in terms of total IMPSE.

In the future, we plan to investigate also other sampling strategies such as taking into account the posterior of the first objective when deciding where to evaluate the second objective, or to learn each objective function's lengthscale and cost on the fly.

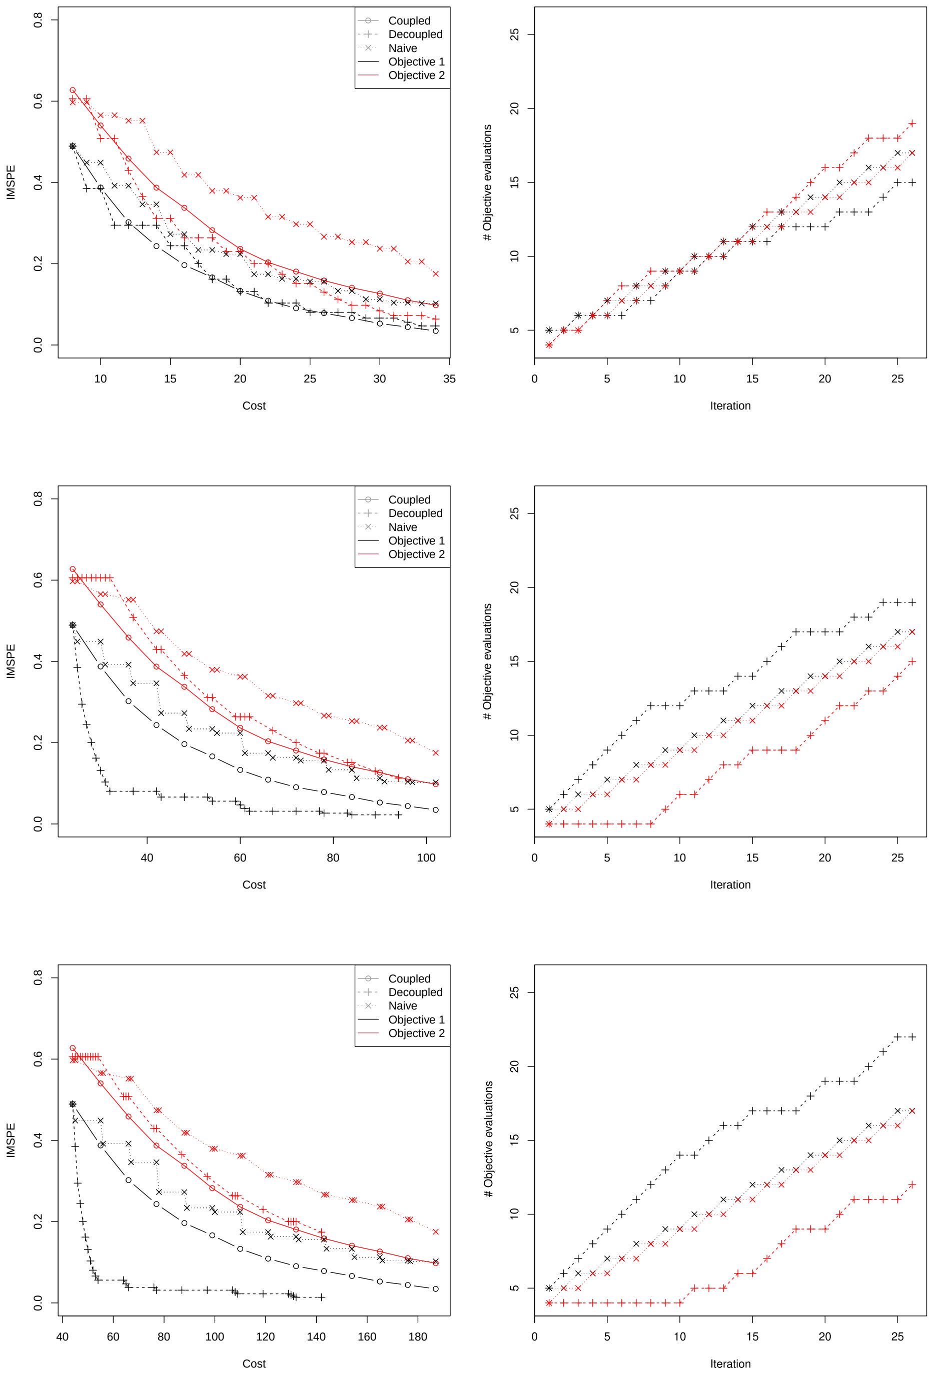

Figure 5 Left: IMSPE vs. cost for the various strategies. Right: objectives evaluated per iteration. Top: cost is equal for both objectives, Middle: cost of f_2 is 5 times greater, Bottom: cost of f_2 is 10 times greater.

4.1.5 Acknowledgements

This report benefited from wider discussions within the "Surrogates" working group of the Dagstuhl Seminar 23361 Multiobjective Optimization on a Budget. This group's members included Thomas Bäck, Mickaël Binois, Jürgen Branke, Jonathan Fieldsend, Ekhine Irurozki, Pascal Kerschke, Boris Naujoks, Robin Purshouse, Tea Tusar, Vanessa Volz, Hao Wang and Kaifeng Yang.

References

1 Richard Allmendinger and Joshua Knowles, Heterogeneous Objectives: State-of-the-Art and Future Research, *arXiv preprint arXiv:2103.15546*, 2, 2021
2 Raul Astudillo and Peter Frazier, Multi-attribute Bayesian optimization under utility uncertainty In: *Proceedings of the NIPS Workshop on Bayesian Optimization*, 172, 2017.
3 Jack M. Buckingham, Sebastian Rojas Gonzalez and Juergen Branke, Bayesian Optimization of Multiple Objectives with Different Latencies, *arXiv preprint arXiv:2302.01310*, 2023.
4 Mickaël Binois, David Ginsbourger and Olivier Roustant, Quantifying uncertainty on Pareto fronts with Gaussian process conditional simulations, *European Journal of Operational Research*, 243(2), 386–394, 2015.
5 Eduardo C Garrido-Merchán and Daniel Hernández-Lobato, Predictive Entropy Search for Multi-objective Bayesian Optimization with Constraints, *Neurocomputing*, 361, 50–68, 2019.
6 Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah and Ryan Adams, Predictive entropy search for multi-objective Bayesian optimization, In: *International Conference on Machine Learning*, 1492–1501, PMLR, 2016.
7 Md Shahriar Iqbal, Jianhai Su, Lars Kotthoff and Pooyan Jamshidi, FlexiBO: A Decoupled Cost-Aware Multi-Objective Optimization Approach for Deep Neural Networks, *Journal of Artificial Intelligence Research*, 77, 645–682, 2023.
8 Nasrulloh Loka, Ivo Couckuyt, Federico Garbuglia, Domenico Spina, Inneke Van Nieuwenhuyse and Tom Dhaene, Bi-objective Bayesian optimization of engineering problems with cheap and expensive cost functions, *Engineering with Computers*, 1, 2022.
9 M. D. Mckay, R. J. Beckman and W. J. Conover, A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code, *Technometrics*, 42(1), 55–61, 2000.
10 Harald Niederreiter, *Random number generation and quasi-Monte Carlo methods*, SIAM, 1992.
11 Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara and Masayuki Karasuyama, Multi-objective Bayesian optimization using Pareto-frontier entropy, In: *International Conference on Machine Learning*, 9279–9288, PMLR, 2020.

4.2 Hypervolume-Indicator-Based Evolutionary Algorithms on a Budget

Jürgen Branke (University of Warwick, GB), Kerstin Dächert (HTW Dresden, DE), Andrzej Jaszkiewicz (Poznan University of Technology, PL), Kathrin Klamroth (Universität Wuppertal, DE)

4.2.1 Motivation

Indicator-based evolutionary algorithms are among the most powerful multi-objective algorithms, in particular when using hypervolume (HV) contribution as indicator. They are not really suitable for many-objective problems, as the computational cost for computing HV

14 Vanessa Volz, Boris Naujoks, Pascal Kerschke, and Tea Tušar. Single- and multi-objective game-benchmark for evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2019, Prague, Czech Republic, July 13-17, 2019*, pages 647–655. ACM, 2019.

5 Seminar Schedule

Monday, September 4, 2023

09:00–10:30 Welcome Session

- Welcome and introduction
- Short presentation of all participants

Coffee Break

11:00–12:00 Seminar Overview

- Seminar scope
- Offers-and-needs market

Lunch

14:00–15:30 Modelling to Address Budget Constraints

- Juliane Mueller: *Surrogate model guided optimization of expensive black-box multiobjective problems*
- Andrea Raith: *Problem decomposition in biobjective optimization*

Coffee Break

16:00–16:30 Working Group Formation

16:30–18:00 Working Groups

Dinner

Tuesday, September 5, 2023

09:00–09:30 Reporting from Small Working Groups Chair: Matthias Ehrgott

09:30–10:30 Small Working Groups

Coffee Break

11:00–12:00 Heuristic Optimization and Human Involvement Chair: Jürgen Branke

- Thomas Bäck: *How to help end users when the budget is limited?*
- Robin Purshouse: *Towards decision analytic workflows for real-world problems: Simulation model calibration and multi-objective optimization on a shared evaluation budget*
- Benjamin Doerr: *Runtime analysis for the NSGA-II*
- Kaisa Miettinen: *Perspectives to dealing with computationally expensive multiobjective optimization problems*

Lunch

14:00–15:30 Small Working Groups

Coffee Break

16:00–17:00 Small Working Groups

17:00–18:00 Reporting from Small Working Groups and General Discussion
Dinner

Wednesday, September 6, 2023

09:00–10:30 Small Working Groups
Coffee Break

11:00–12:00 Approximation and Exact Methods Chair: Karl-Heinz Küfer

- Kathrin Klamroth: *Objective space methods: Pareto front approximations on a budget*
- Michael Stiglmayr: *Multi-objective branch-and-bound on a budget*
- Frank Neumann: *Fast Pareto optimization using sliding window selection*
- Alma Rahat: *Efficient approximation of expected hypervolume improvement using Gauss-Hermite quadrature*

12:00-12:15 Group Photo Outside
Lunch

13:30–15:30 Hiking Trip
Coffee Break

16:00–18:00 Small Working Groups
Dinner

Thursday, September 7, 2023

09:00–10:00 Reporting from Small Working Groups Chair: Kaisa Miettinen
Coffee Break

10:30–12:00 Small Working Groups
Lunch

14:00–15:30 Small Working Groups
Coffee Break

16:00–16:30 Graphical and probabilistic approaches Chair: Boris Naujoks

- Ralph Steuer: *A visualization-aided approach for solving tri-criterion portfolio problems*
- Hao Wang and Kaifeng Yang: *Probability of "improvement" in multi-objective Bayesian optimization*

16:30–17:00 Announcements

17:00–18:00 Joint session with DS22362
Dinner

20:00–23:00 Informal Get Together (BYOB, meet in the cafeteria)

Friday, September 8, 2023

09:00–10:30 Final Reporting from Working Groups
Coffee Break

11:00–12:00 Closing Session

6 Topics of interest for participants for next Dagstuhl Seminar

In the closing session on Friday, the participants reflected upon their experience and presented their ideas on a potential future seminar that would leverage the progress made during the current one. During this discussion, some topics appeared to center around "Artificial Intelligence (AI)". A two-way perspective was suggested: AI for multiobjective optimization and multiobjective optimization for AI. Another suggestion was to focus on the "gap" between the industrial and the academic practice of multiobjective optimization. This suggestion was well-received by both industrial and academic participants of the seminar as the focus during the week was on a "budget" that might also mean decision maker's limitations. Focusing on how the theoretical and methodological achievements on the academic front can be made more accessible to practitioners in industry may be a future direction to pursue. This direction will also possibly require placing more emphasis on modelling, handling the noise, errors and uncertainties in the process. The organizers will use these suggestions as the basis for their discussion about possible topics for the next edition of this seminar series and for the preparation of a proposal for a continuation of the series.

7 Changes in the seminar organization body

As part of a continuing effort to renew the organizing board of this series of Dagstuhl Seminars, Margaret Wiecek steps down from the team of organizers, a role that she has held for three terms of office. On behalf of all the participants of the seminar, Richard Allmendinger, Carlos Fonseca and Serpil Sayin would like to express appreciation to Margaret for her contributions and leadership that have been fundamental for the series success.

We are pleased to announce that our esteemed colleague and a multiple-times Dagstuhl attendee Susan Hunter has agreed to serve as a co-organizer for future editions of this Dagstuhl Seminar series on Multiobjective Optimization. We look forward to collaborating with her in the near future.

Participants

- Thomas Bäck
Leiden University, NL
- Mickaël Binois
INRIA – Sophia Antipolis, FR
- Fritz Bökler
Universität Osnabrück, DE
- Jürgen Branke
University of Warwick, GB
- Dimo Brockhoff
INRIA Saclay – Palaiseau, FR
- Tinkle Chugh
University of Exeter, GB
- Kerstin Dächert
HTW Dresden, DE
- Benjamin Doerr
Ecole Polytechnique – Palaiseau, FR
- Matthias Ehrgott
Lancaster University, GB
- Gabriele Eichfelder
TU Ilmenau, DE
- Jonathan Fieldsend
University of Exeter, GB
- Carlos M. Fonseca
University of Coimbra, PT
- Susan R. Hunter
Purdue University, US
- Ekhine Irurozki
Telecom Paris, FR
- Hisao Ishibuchi
Southern Univ. of Science and Technology – Shenzen, CN
- Andrzej Jaszkiewicz
Poznan University of Technology, PL
- Pascal Kerschke
TU Dresden, DE
- Kathrin Klamroth
Universität Wuppertal, DE
- Joshua D. Knowles
Schlumberger Cambridge Research, GB
- Karl Heinz Küfer
Fraunhofer ITWM – Kaiserslautern, DE
- Arnaud Liefooghe
University of Lille, FR
- Kaisa Miettinen
University of Jyväskylä, FI
- Juliane Mueller
NREL – Golden, US
- Boris Naujoks
TH Köln, DE
- Aneta Neumann
University of Adelaide, AU
- Frank Neumann
University of Adelaide, AU
- Markus Olhofer
HONDA Research Institute Europe GmbH – Offenbach, DE
- Luís Paquete
University of Coimbra, PT
- Robin Purshouse
University of Sheffield, GB
- Alma Rahat
Swansea University, GB
- Andrea Raith
University of Auckland, NZ
- Enrico Rigoni
ESTECO SpA – Trieste, IT
- Stefan Ruzika
RPTU – Kaiserslautern, DE
- Serpil Sayin
Koc University – Istanbul, TR
- Anita Schöbel
Fraunhofer ITWM – Kaiserslautern, DE
- Britta Schulze
Universität Wuppertal, DE
- Ralph E. Steuer
University of Georgia, US
- Michael Stiglmayr
Universität Wuppertal, DE
- Tea Tusar
Jozef Stefan Institute – Ljubljana, SI
- Daniel Vanderpooten
University Paris-Dauphine, FR
- Vanessa Volz
modl.ai – Copenhagen, DK
- Hao Wang
Leiden University, NL
- Margaret M. Wiecek
Clemson University, US
- Kaifeng Yang
Univ. of Applied Sciences – Hagenberg, AT