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

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.

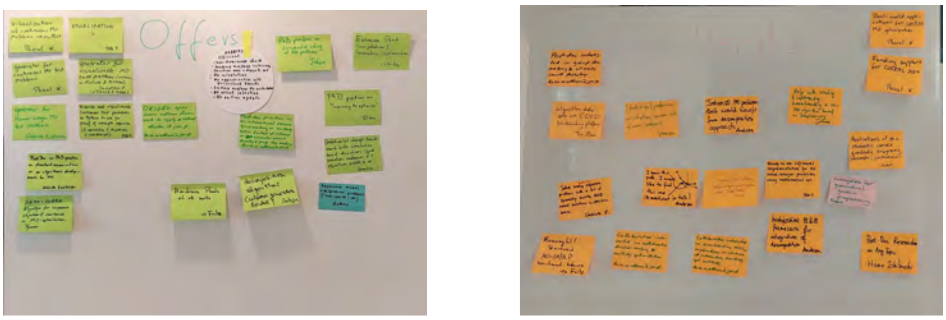


**Figure 1** Offers and needs market.

## 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

## 4.8 Exploring correlations in multi-objective optimization

*Tea Tusar (Jozef Stefan Institute – Ljubljana, SI), Dimo Brockhoff (INRIA Saclay – Palaiseau, FR), Jonathan Fieldsend (University of Exeter, GB), Boris Naujoks (TH Köln, DE), Alma Rahat (Swansea University, GB), Vanessa Volz (modl.ai – Copenhagen, DK), Hao Wang (Leiden University, NL), Kaifeng Yang (Univ. of Applied Sciences – Hagenberg, AT)*



### 4.8.1 Introduction

A frequent assumption in evolutionary computation is that all function evaluations take the same amount of time. However, this rarely holds for real-world optimization problems, especially those that rely on simulations for evaluating solutions. There, the evaluation time can differ for different objectives as well as for different solutions.

The case where evaluation time depends on objectives has already been explored in a previous Dagstuhl Seminar [7]. This typically occurs in problems where some objectives can be computed with a closed-form expression while others require lengthy simulations. Various strategies for handling objectives with heterogeneous evaluation times are reviewed in [1].

During this seminar, we focused on the second case, in which the evaluation time depends on solutions. Specifically, we wanted to explore whether the correlation between objectives and their evaluation times can be modeled and exploited to save expensive function evaluations.

### 4.8.2 Motivation from real-world applications

In some real-world problems, the relation between solution properties and evaluation times is rather straightforward. For example, in the tunnel alignment problem [9], where a solution represents a tunnel trajectory, the computational expense of assessing tunnel objectives and constraints is proportional to the length of the tunnel – a longer tunnel will generally take longer to evaluate. Similarly holds for neural architecture search [3], where a solution defines the architecture of a neural network whose training time is strongly positively correlated with its size.

However, there are also other kinds of real-world problems where such a relation is hard to find. Consider the airfoil optimization problem [11], where computational fluid dynamics is used in solution evaluation, and the electrical motor design problem [13], which relies on electromagnetic field simulations. In both cases, evaluation times vary among solutions, but a clear correlation between solution characteristics and evaluation duration has not been discovered.

Another source of solution-dependent evaluation times is the presence of hidden constraints. For instance, the MarioGAN optimization problem [14] involves generating Mario game levels, which are assessed through playthrough simulations with artificial intelligence players. If a generated level cannot be solved (that is, Mario cannot reach the level end), the simulation would continue endlessly unless terminated. The distance in the search space between feasible solutions that are relatively quick to evaluate and infeasible solutions whose evaluation takes a long time can be very small in such cases.

These examples show that the correlation between objective quality and its evaluation time depends on the problem and the solutions. We can model it by considering the evaluation time as an additional independent objective to be minimized.

### 4.8.3 Visualization of correlations

We use search space visualizations to gain a better understanding of the correlations between objectives. The correlation for each pair of objectives is estimated in different regions of the search space using the Pearson correlation coefficient for a small (local) sample of the search space. The Pearson correlation coefficient measures the linear correlation between two samples' objectives and takes a value between $-1$ (perfect linear anti-correlation) and 1 (perfect linear correlation). A 0 value implies that there is no linear dependency between the objectives. The Pearson correlation coefficient is invariant when the two objectives are shifted and/or scaled.

#### 4.8.3.1 Experimental setup

For demonstration purposes, we choose some continuous test problems with 2-D search spaces that are straightforward to visualize. They have either two, three or five objectives and various characteristics (more details below). We assume minimization of their $m$ objectives.

The 2-D problem search space is discretized into a grid of $501 \times 501$ points. For each grid point $x = (x_1, x_2)$, the correlation between two objectives is computed with the Pearson correlation coefficient as follows. First, $p$ equidistant points are created on the circle with radius $10^{-6}$ centered at $(x_1, x_2)$ with one point placed at position $(x_1 + 10^{-6}, x_2)$, see Figure 14. Next, the $p$ points are evaluated, i.e. $m$ objective values are computed for each of them. Finally, the correlation between each pair of objectives at $x$ is estimated with the corresponding Pearson correlation coefficient for the set of $p$ points. Note that the $p$ points could have been constructed also in some other way. We opted for this deterministic approach to minimize the disturbances caused by a stochastic choice of point placement. In all experiments, the number of points $p$ was set to 100.

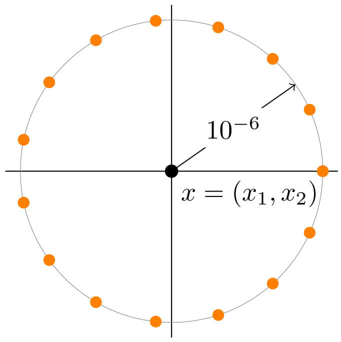


**Figure 14** The grid point $x = (x_1, x_2)$ and the $p$ points (shown in orange) used in the computation of the Pearson correlation coefficient (here, $p = 15$).

#### 4.8.3.2 Problems with two objectives

First, we wish to explore the simplest case of two objectives. For this, we select six bi-objective problems from the `bbob-biobj` suite of benchmark problems [2]. They are constructed by combining two single-objective functions from the `bbob` suite [8]. Figure 15 shows the visualization of correlations between the two objectives for each of the six problems.

The double sphere problem $F_1 = (f_1, f_1)$, where $f_1$ is the `bbob` sphere function, is a unimodal problem with a known Pareto set – the line segment connecting the two single-objective optima. We can see from the correlation plot in Figure 15a the expected outcome – close to the Pareto set, the objectives are anti-correlated (red hues), while further away they are correlated (blue hues).

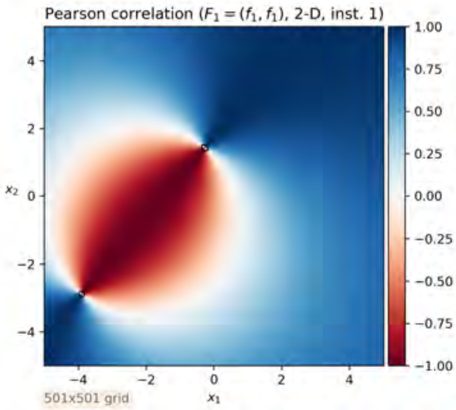


**(a)** The double sphere problem $F_1$.

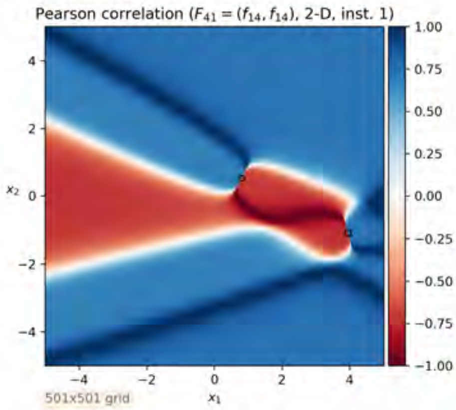


**(b)** The double different powers problem $F_{41}$.

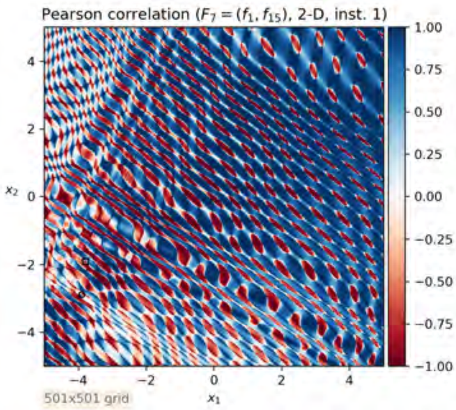


**(c)** The sphere-Rastrigin problem $F_7$.

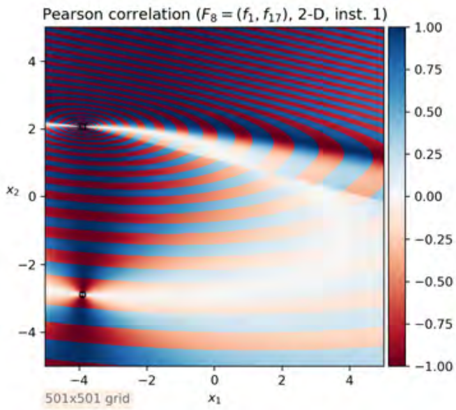


**(d)** The sphere-Schaffer problem $F_8$.

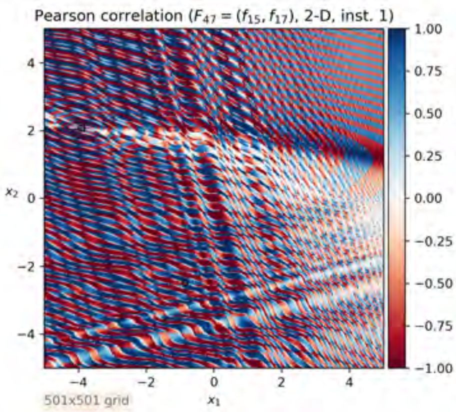


**(e)** The Rastrigin-Schaffer problem $F_{47}$.

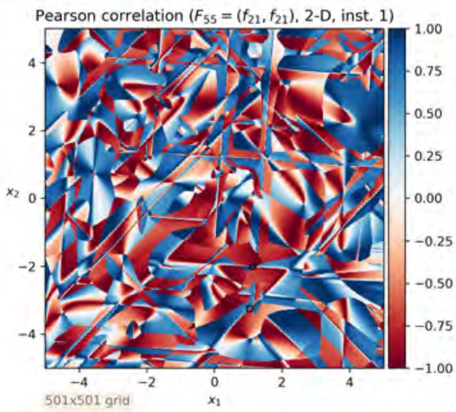


**(f)** The double Gallagher problem $F_{55}$.

**Figure 15** Person correlation coefficient for some chosen 2-D `bbob-biobj` problems (these and other plots for `bbob-biobj` problems will be made available at `https://numbbo.github.io/bbob-biobj/vis/`). Blue hues denote positive correlations, red hues negative ones and white indicates no correlation.

In the problem $F_{41} = (f_{14}, f_{14})$, both objectives are unimodal as well, but they correspond to the bbob sum of different powers function $f_{14}$, which is non-separable and ill-conditioned. Figure 15b shows that in this case, the objectives are anti-correlated also far away from the Pareto set.

The next two problems are a combination of a unimodal objective (the bbob sphere function $f_1$) and a highly multimodal one. In the problem $F_7 = (f_1, f_{15})$, this is the bbob Rastrigin function $f_{15}$, while in the problem $F_8 = (f_1, f_{17})$, it is the bbob Schaffer F7 function $f_{17}$ with condition number 10. In both instances, visualized in Figures 15c and 15d, the resulting bi-objective problems have multiple disconnected regions of the search space where the objectives are anti-correlated.

Finally, in the last two selected problems, both objectives are highly multimodal. The problem $F_{47} = (f_{15}, f_{17})$ combines the bbob Rastrigin function $f_{15}$ with the bbob Schaffer F7 function $f_{17}$ with condition number 10 and the problem $F_{55} = (f_{21}, f_{21})$ two bbob Gallagher's Gaussian functions $f_{21}$ with 101 median peaks. We can see from the correlation plots in Figures 15e and 15f the high number of disconnected regions of anti-correlated objectives.

These examples challenge some of our preexisting notions about the correlation between objectives. In particular, they show that it is closely connected to the problem multimodality – understandably, given that the correlation between two objectives equals $-1$ at any locally optimal set. In fact, the notion of a globally (i.e., Pareto) optimal set is inconsequential for correlation values. It is therefore rather meaningless to discuss correlations between objectives without taking into account their multimodality. We also see that the Pearson correlation coefficient values are themselves positively correlated with the length of the normalized bi-objective gradient as defined in [10] and visualized in [2].

#### 4.8.3.3 Problems with three objectives

The Pearson correlation coefficient is defined only for two objectives. When the objectives are three (or more), we can compute all their pairwise correlations. We wish to visualize their minimal values to emphasize parts of the search space with the highest anti-correlation as they are locally optimal.

Exemplary three-objective problems are again constructed by combining bbob functions – now three. This time, we chose the triple sphere problem, the sphere-Rastrigin-Schaffer problem and the triple Gallagher problem. See Figure 16 for their visualizations. For each problem we show on the left hand side the pairwise correlations for objectives 1 and 2, objectives 1 and 3 and objectives 2 and 3 as well as their mean. On the right hand side, their minimum is presented.

The Pareto set of the triple sphere problem is the triangle spanned by the three single-objective optima. From Figure 16a we see that its minimal pairwise Pearson correlation coefficient equals $-1$ only at the edges of this triangle, not in its interior. This shows that, unlike in the bi-objective case, one cannot rely on pairwise correlations alone to infer local optimality of a solution in case of more than two objectives. A procedure similar to the one from [12] should be tried to amend this issue.

Further examples show the minimal pairwise correlation for the sphere-Rastrigin-Schaffer problem (Figure 16b) and the triple Gallagher problem (Figure 16c). Both are highly multimodal, resulting in many disconnected regions with anti-correlated pairs of objectives (red hues).

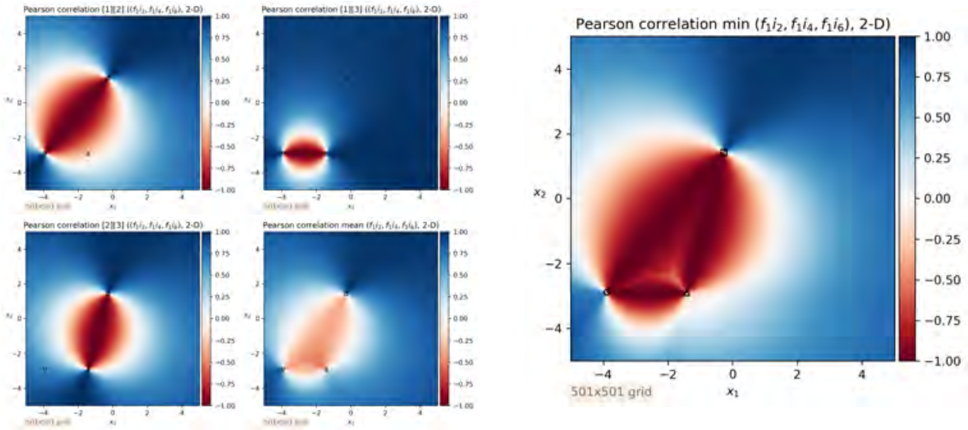


**(a)** The triple sphere problem.

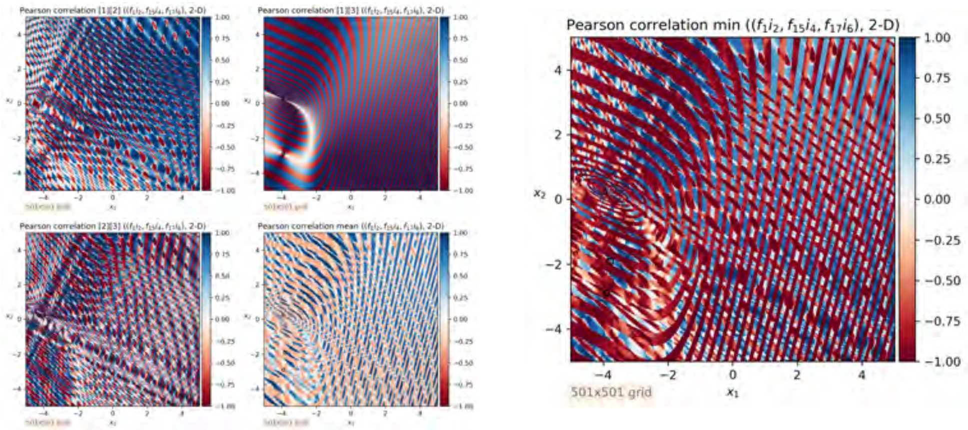


**(b)** The sphere-Rastrigin-Schaffer problem.

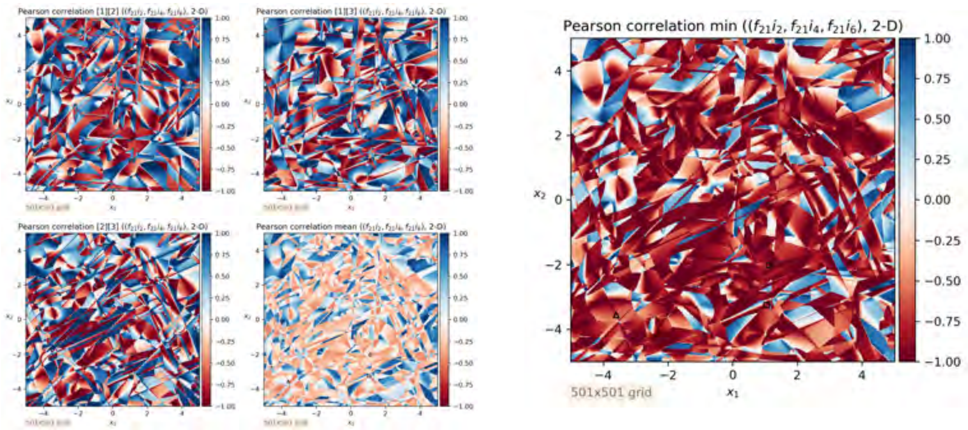


**(c)** The triple Gallagher problem.

**Figure 16** Visualization of correlations for three three-objective problems. Smaller plots from top to bottom, left to right: pairwise Pearson correlation coefficients for objectives 1 and 2, 1 and 3 and 2 and 3, and their mean. Larger plot: minimum value of the pairwise correlation coefficients.

#### 4.8.3.4 Problems with five objectives

We next consider a couple of planar problems with five objectives, using the distance-based multi-objective point problem (DBMOPP) generator [5]. This generator allows us to create problem instances which natively live in 2-D (or map to 2-D), which can have an arbitrary number of objectives and can exhibit a range of other problem properties.

We first generate a box-constrained instance with a single spatially contiguous Pareto set (shown in red in Figure 17a) and seven other regions which generate local fronts of the same shape, but which are dominated (shown in green in Figure 17a). Figure 17b shows the corresponding *dominance landscape* [4]. Black regions in this figure show locations which are not dominated by any immediate neighbor (dominance neutral regions). Gray regions in contrast denote locations which have at least one dominating neighbor, but where all point-based dominance hill-climbs (by moving to an adjacent dominating neighbor) lead to the same dominance neutral region – different shades of gray are used to distinguish these different *basins*. White regions signify where point-based dominance hill-climbs lead to multiple different dominance neutral regions (effectively multi-objective saddle-points), depending on which chain of dominating neighbors one follows. Figure 17c shows the dominance ratio [6] landscape for the problem instance. In this plot, the value at a location denotes the proportion of the entire domain which *weakly dominates* it (i.e. dominates or is equal to it). That is, a value of 0.0 will indicate a location is Pareto optimal, whereas a value of 0.2 indicates that 20% of the domain relates to locations with equal or better performance on all criteria. Pearson correlation plots are shown in Figures 17d–17f. For this problem we can see the eight distinct local optima regions clearly in the Dominance ratio plot, with the induced dominance neutral plateaus between these regions additionally identifiable in the dominance landscape and correlation plots.

The second example shown in Figure 18a has a single spatially contiguous Pareto set region (red), 3 dominance resistance regions (blue), 3 local fronts regions (green) and 30% of the decision space is designed as being flat under the objectives (cyan). The corresponding dominance landscape is shown in 18b, and the dominance ratio landscape in 18c. Pearson correlation plots are shown in Figures 18d–18f. The impact of the flat objective regions is clear across the plots, and all views of the landscape are considerably more cluttered due to the interactions of the various problem features.

### 4.8.4 Conclusions

We recognized that evaluation times can differ among solutions of expensive real-world problems. We were therefore interested in exploring whether the correlation between objectives and their evaluation times can be used to save time-consuming function evaluations. A deeper look into the properties of some real-world applications has shown that a general model for such a correlation is hard to find. Therefore, the evaluation time was regarded as an additional objective to be minimized.

Next, we researched the correlation between objectives, estimating it with the Pearson correlation coefficient. To gain a better understanding of the distribution of its values in the search space, we visualized them for a number of test problems with two variables and two, three and five objectives. The visualizations have shown that some of our intuition about the correlation between objectives was wrong. For example, we could find unimodal problems with anti-correlated objectives not only close to the Pareto set, but also far away from it. Visualizations of multimodal problems have proven that many distinct anti-correlated regions can be located throughout the search space, surrounded by regions with correlated

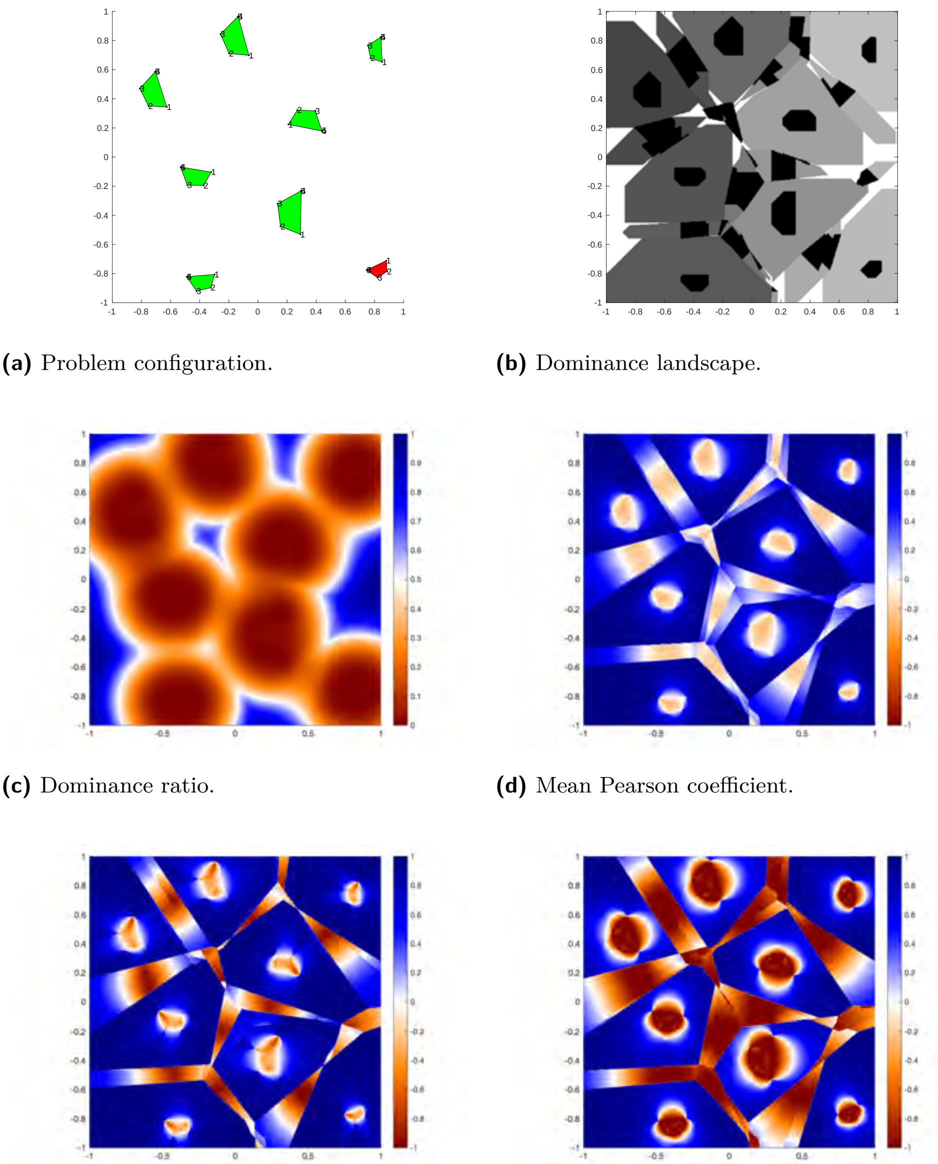

**(a)** Problem configuration. **(b)** Dominance landscape.

**(c)** Dominance ratio. **(d)** Mean Pearson coefficient.

**(e)** Median Pearson coefficient. **(f)** Min Pearson coefficient.

**Figure 17** Problem plots and Person correlation values for a 5-objective 2-D `DBMOPP` instance. In the correlation plots blue hues denote positive correlations, red hues negative ones and white indicates no correlation.

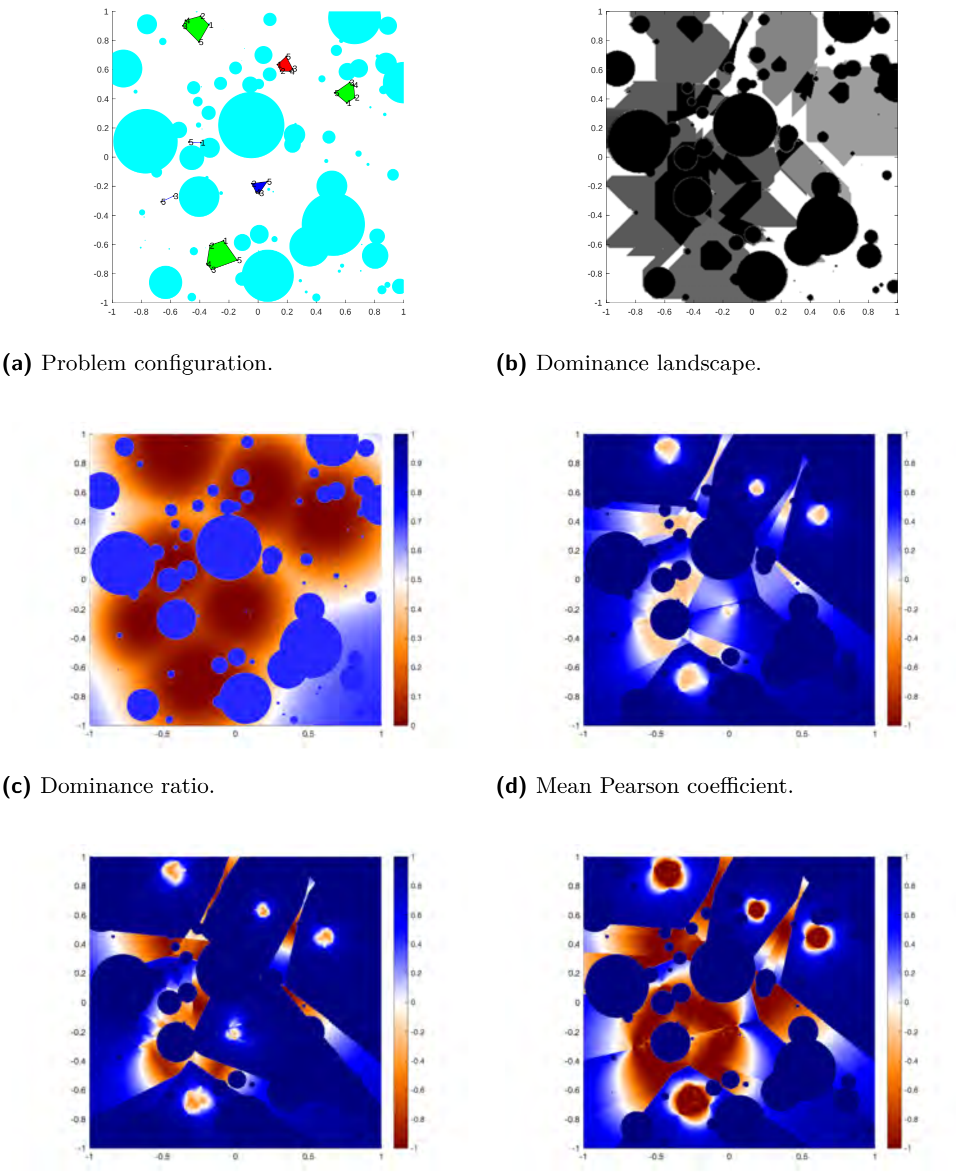

**(a)** Problem configuration.

**(b)** Dominance landscape.

**(c)** Dominance ratio.

**(d)** Mean Pearson coefficient.

**(e)** Median Pearson coefficient.

**(f)** Min Pearson coefficient.

**Figure 18** Problem plots and Person correlation values for a more complex 5-objective 2-D `DBMOPP` instance. In the correlation plots blue hues denote positive correlations, red hues negative ones and white indicates no correlation.

objectives. In fact, the visualizations have demonstrated that correlation is closely tied to the problem multimodality and has a nonlinear monotonous relation with the length of the bi-objective gradient. Finally, while pairwise anti-correlations between objectives correspond to the locally optimal solutions for problems with two objectives, this is no longer the case when the number of objectives is three or more.

### References


1 Richard Allmendinger and Joshua D. Knowles. Heterogeneous objectives: State-of-the-art and future research. *CoRR*, abs/2103.15546, 2021.
2 Dimo Brockhoff, Anne Auger, Nikolaus Hansen, and Tea Tušar. Using well-understood single-objective functions in multiobjective black-box optimization test suites. *Evolutionary Computation*, 30(2):165–193, 2022.
3 Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.
4 Jonathan E. Fieldsend, Tinkle Chugh, Richard Allmendinger, and Kaisa Miettinen. A feature rich distance-based many-objective visualisable test problem generator. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '19, page 541–549, New York, NY, USA, 2019. Association for Computing Machinery.
5 Jonathan E. Fieldsend, Tinkle Chugh, Richard Allmendinger, and Kaisa Miettinen. A visualizable test problem generator for many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 26(1):1–11, 2022.
6 Carlos Manuel Mira da Fonseca. *Multiobjective genetic algorithms with application to control engineering problems.* PhD thesis, University of Sheffield, 1995.
7 Salvatore Greco, Kathrin Klamroth, Joshua D. Knowles, and Günter Rudolph. Understanding complexity in multiobjective optimization (Dagstuhl Seminar 15031). *Dagstuhl Reports*, 5(1):96–163, 2015.
8 Nikolaus Hansen, Steffen Finck, Raymond Ros, and Anne Auger. Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Research Report RR-6829, INRIA, 2009.
9 Miha Hafner, David Rajšter, Marko Žibert, Tea Tušar, Bernard Ženko, Martin Žnidaršič, Flavio Fuart, and Daniel Vladušič. Artificial intelligence support for tunnel design in urban areas. In *Tunnels and Underground Cities: Engineering and Innovation meet Archaeology, Architecture and Art- Proceedings of the WTC 2019 ITA-AITES World Tunnel Congress*, pages 2196–2205, 2019.
10 Pascal Kerschke and Christian Grimme. An expedition to multimodal multi-objective optimization landscapes. In *Evolutionary Multi-Criterion Optimization – 9th International Conference, EMO 2017, Münster, Germany, March 19-22, 2017, Proceedings*, volume 10173 of *Lecture Notes in Computer Science*, pages 329–343. Springer, 2017.
11 Boris Naujoks, Werner Haase, Jörg Ziegenhirt, and Thomas Bäck. Multi objective airfoil design using single parent populations. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference, New York, USA, 9-13 July 2002*, pages 1156–1163. Morgan Kaufmann, 2002.
12 Lennart Schäpermeier, Christian Grimme, and Pascal Kerschke. To boldly show what no one has seen before: A dashboard for visualizing multi-objective landscapes. In *Evolutionary Multi-Criterion Optimization – 11th International Conference, EMO 2021, Shenzhen, China, March 28-31, 2021, Proceedings*, volume 12654 of *Lecture Notes in Computer Science*, pages 632–644. Springer, 2021.
13 Tea Tušar, Peter Korošec, and Bogdan Filipič. A multi-step evaluation process in electric motor design. In *Slovenian Conference on Artificial Intelligence, Proceedings of the 26rd International Multiconference Information Society – IS 2023*, volume A, pages 48–51. Jožef Stefan Institute, 2023.

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