



A game theoretic perspective on Bayesian multi-objective optimization

Mickael Binois, Abderrahmane Habbal, Victor Picheny

► To cite this version:

Mickael Binois, Abderrahmane Habbal, Victor Picheny. A game theoretic perspective on Bayesian multi-objective optimization. Many-Criteria Optimization and Decision Analysis: State-of-the-Art, Present Challenges, and Future Perspectives, 2023, 978-3-031-25263-1. 10.1007/978-3-031-25263-1_11 . hal-03206174

HAL Id: hal-03206174

<https://inria.hal.science/hal-03206174>

Submitted on 29 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A game theoretic perspective on Bayesian multi-objective optimization

Mickaël Binois* Abderrahmane Habbal* Victor Picheny†

Abstract

This chapter addresses the question of how to efficiently solve many-objective optimization problems in a computationally demanding black-box simulation context. We shall motivate the question by applications in machine learning and engineering, and discuss specific harsh challenges in using classical Pareto approaches when the number of objectives is four or more. Then, we review solutions combining approaches from Bayesian optimization, e.g., with Gaussian processes, and concepts from game theory like Nash equilibria, Kalai-Smorodinsky solutions and detail extensions like Nash-Kalai-Smorodinsky solutions. We finally introduce the corresponding algorithms and provide some illustrating results.

1 Introduction

In the previous Chapter 11, methods for multi-objective (MO) Bayesian optimization (BO) have been introduced. Scaling with many-objectives burst the seams of these existing solutions on several aspects: i) The positive-definiteness requirement for multi-output Gaussian process (GP) covariance matrix severely restricts their modelling ability, letting independent GPs as the default alternative that does not directly exploit possible correlations between objectives (see also Chapter 10); ii) Summing the objectives is a go-to strategy often observed in practice. But picking weights in scalarization methods becomes more arbitrary, with effects that are difficult to apprehend. More generally, specifying preferences for a large number of conflicting objectives can lead to situations where the connection between achievable solutions and the given preferences is not easy to interpret; iii) For Pareto-based infill criteria, the computational cost grows quickly, for instance with hypervolume computations. Plus there are less closed form expressions available than for two or three objectives; iv) The effect of more objectives on the acquisition function landscape is to be studied, but they are expected to be more multimodal due to larger Pareto sets. And these are only a few additional technical problems faced by many-objective optimization (MaOO).

But even the desired outcome is problematic. As the number of objectives increases, the dimension of the Pareto front generally grows accordingly, and representing this latter is more complex. With a limited budget of evaluations, this may not even be achievable: a few hundred design points would not give an accurate representation of a complex ten-dimensional manifold. In addition, even if representing the entire Pareto front was achievable, the decision maker would be left with a very large set of incomparable solutions, with limited practical relevance.

*Inria, Université Côte d’Azur, CNRS, LJAD, UMR 7351, Sophia Antipolis, France, mickael.binois@inria.fr, habbal@unice.fr

†Secondmind, Cambridge, United Kingdom, victor@secondmind.ai

Hence, MaOO is more of an elicitation problem, for which good and principled ways to choose solutions are needed. When the decision maker is involved in the optimization process by giving feedback on pairs of solutions, methods learning this preference information such as [33, 5] are available. More diffuse preference information can also be included via the choice of the reference point for hypervolume computation (nullifying the contributions of equivalent or dominated solutions), interactively modifying ranges for objectives [26] or via an aspiration point serving as target (see, e.g., [20]). Ordering the importance of objectives is another option, exploited e.g., in [1]. With many objectives, possibly of various natures and scales, the burden on the decision maker is heavier, and the effect on the outcome of these preferences is less intuitive.

When the decision is only made *a posteriori*, knee points, Pareto solutions from which improving slightly one objective would deteriorate much more some others, are appealing to decision makers, or in the design of evolutionary algorithms, see e.g., [52]. But their existence and location on the Pareto front may vary, especially in the MaOO case. Other arguments have been directed toward finding a single solution at the “center” of the Pareto front to keep balance between objectives, e.g., by [21, 39]. The rationale is that trade-off solutions at the extremes are not desirable unless prior information has been given. Finding such balanced solutions, without input from the decision maker during the optimization, is the default framework considered in this chapter. The difficulty lies in an appropriate and principled definition of a central solution. A possible approach is to frame the MaOO problem as a game theory one, where each objective is impersonated by a *player*, and the players bargain until they agree on a satisfying solution for all, referred to as *equilibrium*.

Game theory is already popular in a variety of applications such as engineering [15], control [25], multi-agent systems [11] or machine learning [32]. Game theory has also revealed to be a powerful tool in multi-disciplinary optimization, where decentralized strategies for non-cooperative agents takes its full and relevant meaning, as referring to discipline-specific solvers in a parallel, asynchronous and heterogeneous multi-platform setting, see e.g., the pioneering framework developed in [40]. It comes equipped with principled ways of eliciting solutions (the *bargaining axioms* of scale invariance, symmetry, efficiency, and independence of irrelevant alternatives), echoing several desirable properties for improvement functions used in MO BO [45]: reflecting Pareto dominance, no input in the form of external parameters (such as a reference point), or robustness to scale invariance. Indeed, with more objectives of different nature, finding a reasonable common scale for all is much harder and the effects cannot be apprehended.

Compared to the more standard multi-objective BO framework described in the previous chapter, the games framework comes with some specific challenges. The definition of solutions such as Nash equilibrium or Kalai-Smorodinsky solutions can be quite involved, and finding such equilibria generally implies solving several inference tasks at once. For instance, this makes the use of improvement-based approaches difficult. Still, several approaches have emerged using regret or stepwise uncertainty reduction (SUR) frameworks, see e.g., [42, 1] and references therein.

The structure of the chapter is as follows. In Section 2, we briefly review key concepts in game theory such as Nash equilibria. Then in Section 3 we present high-level BO approaches to solve many-objective problems under the Games paradigm. Illustrations are provided in Section 4 before discussing remaining challenges and perspectives in Section 5.

2 Game equilibria to solution elicitation

The standard multi-objective optimization problem (MOP) corresponds to the simultaneous optimization of all objectives:

$$(MOP) \quad \min_{\mathbf{x} \in \mathcal{X}} (f_1(\mathbf{x}), \dots, f_m(\mathbf{x})). \quad (1)$$

Besides considering objectives as individual players' goals, general non-cooperative games, and related solutions, such as Nash ones, do need territory splitting. That is, partitioning of the optimization variables (the input space) among players. Other games and related solution concepts depend on choices of anchor points, as is the case for the Kalai-Smorodinsky (KS) solution, which depends on ideal and disagreement points, as illustrated in Figure 1. We shall introduce in the following the Nash equilibrium concept, then motivated by the generic inefficiency of the latter, move to considering the Kalai-Smorodinsky solution.




Figure 1: Illustration of solutions of two different elicitation problems with two different disagreement points \mathbf{d} : either the Nadir point or a Nash equilibrium, resulting in the KS and NKS points (stars) on the Pareto front \mathcal{P} , respectively. The shaded area shows the feasible objective space.

2.1 Nash games and equilibria

When considering primarily the standard (static, under complete information) Nash equilibrium problem [23], each objective becomes a player's outcome. Compared to a standard MOP (1) where players share the control of the same set of variables (or *action space*), variables are uniquely allocated to a player in the so-called territory splitting. Denote \mathbf{x}_i the variables of player i and \mathcal{X}_i its corresponding action space, where $\mathcal{X} = \prod_i \mathcal{X}_i$. Accordingly, the variable vector \mathbf{x} consists of block components $\mathbf{x}_1, \dots, \mathbf{x}_m$ ($\mathbf{x} = (\mathbf{x}_j)_{1 \leq j \leq m}$). We shall use the convention $f_i(\mathbf{x}) = f_i(\mathbf{x}_i, \mathbf{x}_{-i})$ when we need to emphasize the role of \mathbf{x}_i , where \mathbf{x}_{-i} is the subset of variables controlled by players $j \neq i$.

Definition 2.1. A Nash equilibrium problem (NEP) consists of $m \geq 2$ decision makers (i.e., players), where each player $i \in \{1, \dots, m\}$ wants to solve its optimization problem:

$$(\mathcal{P}_i) \quad \min_{\mathbf{x}_i \in \mathcal{X}_i} f_i(\mathbf{x}), \quad (2)$$

and $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), \dots, f_m(\mathbf{x})] : \mathcal{X} \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ (with $n \geq m$) denotes a vector of cost functions (a.k.a. pay-off or utility functions when maximized), f_i denotes the specific cost function of player i .

Definition 2.2. A Nash equilibrium $\mathbf{x}^* \in \mathcal{X}$ is a strategy such that:

$$(NE) \quad \forall i, 1 \leq i \leq m, \quad \mathbf{x}_i^* = \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} f_i(\mathbf{x}_i, \mathbf{x}_{-i}^*). \quad (3)$$

In other words, when all players have chosen to play a NE, then no single player has incentive to move from his \mathbf{x}_i^* .

The above definition shows some attractive features of NE: it is scale-invariant, does not depend on arbitrary parameters (such as the reference point), and convey a notion of first-order stationarity. On the other hand, it requires partitioning the inputs between players, which may come naturally for some problems, but may be problematic for many MaOO problems. Plus an arbitrary partitioning can have an unknown outcome on the solution. One option proposed in [17] is to define the partitioning based on sensitivity analysis of one main objective. Hence a perspective would be to define an optimal partitioning in the sense that it allocates variables (or a linear combination of variables) to the player it has most influence on.

Besides the need to define a territory splitting, it is important to notice that, generically, Nash equilibria are not efficient, i.e., do not belong to the underlying set of best compromise solutions, the so-called Pareto front, of the objective vector $(f_i(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}}$. Indeed, NE efficiency may happen when one of the players has control on all the optimization variables while the others have control on nothing (a somehow degenerate territory splitting). Still, NE have a couple of advantages: a nice notion of stationarity, they are generally well balanced and, of interest for many objectives, scale invariant.

Remark. When block components $\mathbf{x}_1, \dots, \mathbf{x}_m$ of the variable vector \mathbf{x} overlap due to, e.g., constraints, then one should modify (2) and (3) accordingly, so that $\mathbf{x}_i \in \mathcal{X}_i(\mathbf{x}_{-i})$. In the literature, the Nash equilibrium problem is then referred to as a generalized Nash equilibrium problem (GNEP).

Remark. The NE solution is scale invariant, and more generally invariant under any strictly increasing transformation Ψ_i ($1 \leq i \leq m$) since, in this case, (3) is equivalent to

$$\forall i, 1 \leq i \leq m, \quad \Psi_i(f_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*)) \leq \Psi_i(f_i(\mathbf{x}_i, \mathbf{x}_{-i}^*))$$

Interestingly, it could be shown at least in the convex case that if players share probabilistic control on the optimization variable components (i.e., any player j may control the same component \mathbf{x}_i with probability p_{ij}) then there exists a probability matrix (p_{ij}) such that the associated Nash equilibrium (with an *ad hoc* definition) lies in the Pareto front, see [2] for a sketch of the approach.

From a practical viewpoint, computing NE for games stated in continuous variable settings (as opposed to discrete or finite games, e.g., in vector spaces with Banach or Hilbert structure) can be based on variational analysis, e.g., with the classical fixed-point algorithms to solve NEPs [48, 6, 35]. A modified notion of Karush-Kuhn-Tucker (KKT) points, adapted to (generalized) NE, is developed by [30], to propose a dedicated augmented Lagrangian method. Yet these methods require too many evaluations to tackle directly expensive black-boxes, requiring specific BO algorithms that are presented in Section 3.

2.2 The Kalai-Smorodinsky solution

The Kalai-Smorodinsky solution was first proposed by Kalai and Smorodinsky in 1975 as an alternative to the Nash bargaining solution in cooperative bargaining. Differently from Nash equilibria, no partitioning of the decision variable among players is required by the KS solution concept.

The problem is as follows: starting from a *disagreement* or *status quo* point \mathbf{d} in the objective space, the players aim at maximizing their own benefit while moving from \mathbf{d} toward the Pareto front (i.e., the efficiency set). The KS solution is of egalitarian inspiration [14] and states that the selected efficient solution should yield equal benefit ratio to all the players. Indeed, given the utopia (or ideal, or shadow) point $\mathbf{u} \in \mathbb{R}^m$ defined by

$$u_i = \min_{\mathbf{x} \in \mathcal{X}} f_i(\mathbf{x}),$$

selecting any compromise solution $\mathbf{y} = [f_1(\mathbf{x}), \dots, f_m(\mathbf{x})]$ would yield, for objective i , a benefit ratio

$$r_i(\mathbf{x}) = \frac{d_i - y_i}{d_i - u_i}.$$

Notice that the benefit from staying at \mathbf{d} is zero, while it is maximal for the generically unfeasible point \mathbf{u} . The KS solution is the Pareto optimal choice $\mathbf{y}^* = [f_1(\mathbf{x}^*), \dots, f_m(\mathbf{x}^*)]$ for which all the benefit ratios $r_i(\mathbf{x})$ are equal:

$$\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \text{ s.t. } r_1(\mathbf{x}) = \dots = r_m(\mathbf{x})\} \cap \{\mathbf{x} \in \mathcal{X} \text{ s.t. } \mathbf{f}(\mathbf{x}) \in \mathcal{P}\}.$$

Geometrically, \mathbf{y}^* is the intersection point of the Pareto front and the line (\mathbf{d}, \mathbf{u}) (see Figure 1), which may be empty. For $m = 2$, \mathbf{y}^* exists if the Pareto front is continuous. For the general case, we use here the extension of the KS solution proposed by [28] under the name *efficient maxmin solution*. Since the intersection with the (\mathbf{d}, \mathbf{u}) line might not be feasible, there is a necessary trade-off between Pareto optimality and centrality. The efficient maxmin solution is defined as the Pareto-optimal solution that maximizes the smallest benefit ratio among players, that is:

$$\mathbf{x}^{**} \in \arg \max_{\mathbf{y} \in \mathcal{P}} \min_{1 \leq i \leq m} r_i(\mathbf{x}).$$

It is straightforward that when the intersection is feasible, then \mathbf{y}^* and $\mathbf{y}^{**} := \mathbf{f}(\mathbf{x}^{**})$ coincide. Figure 1 shows \mathbf{y}^{**} in the situation when the feasible space is nonconvex. \mathbf{y}^{**} is always on the Pareto front (hence not necessarily on the (\mathbf{d}, \mathbf{u}) line). In the following, we refer indifferently to \mathbf{y}^* (if it exists) and \mathbf{y}^{**} as the KS solution. Note that this definition also extends to discrete Pareto sets.

For $m = 2$, the (non-extended) KS solution can be axiomatically characterized as the unique solution fulfilling all the bargaining solution axioms, which are: Pareto optimality, symmetry, affine invariance, and restricted monotonicity [29]. For $m \geq 3$, there is no axiomatic setting, the KS solution being required to fulfill Pareto optimality, affine invariance, and equity in benefit ratio. Still, KS is particularly attractive in a many-objective context since it scales naturally to a large number of objectives and returns a single solution, avoiding the difficulty of exploring and approximating large m -dimensional Pareto fronts—especially with a limited number of observations.

Closely related to the Kalai and Smorodinsky solution is the reference point methodology, developed by [51]. The reference point methodology uses achievement functions, which refer to aspiration and reservation references. Ideal and nadir points are particular instances of such reference points, respectively, and achievement functions play the same role as the benefit ratio of our present setting. Moreover, the notion of neutral compromise and max-min

(of achievement functions) were introduced by Wierzbicki in the cited reference, yielding a framework objectively close to the one of Kalai and Smorodinsky. One important difference between the reference point and KS approaches is that the latter relies on a game theoretic axiomatic construction, in view of palliating the Nash equilibrium inefficiency.

2.3 Disagreement point choice

Clearly, the KS solution strongly depends on the choice of the disagreement point \mathbf{d} . A standard choice is the nadir point N given by $N_i = \max_{\mathbf{x} \in \text{Pareto set}} f_{(i)}(\mathbf{x})$. Some authors introduced alternative definitions, called extended KS, to alleviate the critical dependence on the choice of \mathbf{d} [10], for instance by taking as disagreement point the Nash equilibrium arising from a previously played non-cooperative game. But such a choice would need a pre-bargaining split of the decision variable vector \mathbf{x} among the m players. When some relevant territory splitting is agreed on, then Nash yields KS an interesting and non arbitrary disagreement point, and in return, KS yields the Nash game an interesting and non arbitrary efficient solution, starting from the non efficient equilibrium. The resulting Nash-Kalai-Smorodinsky solution is denoted NKS.

The nadir point is very useful to rescale the Pareto front in case objectives are of different nature or scale, see e.g., [7]. As such, it makes a natural disagreement solution. Still, finding the nadir point is a complex task, especially under a limited budget, as it involves exploring regions good for a few objectives with worse values on others. The many objective context makes this issue more prominent. It is related to the problem of Pareto resistance, i.e., regions of the space with good values on a few objectives but not Pareto optimal (see, e.g., [19] and references therein). Unless some objectives are more important than others, the corresponding extremal regions of the Pareto front are of little interest for selecting a single good solution.

A simpler disagreement point is via the *pseudo-nadir* point, defined as the worst objective solution over designs achieving the minimum for one objective: $\tilde{N}_i = \max_{\mathbf{x} \in \{\mathbf{x}^{(1)*}, \dots, \mathbf{x}^{(m)*}\}} f_i(\mathbf{x})$, where $\mathbf{x}^{(j)*} \in \arg \min_{\mathbf{x} \in \mathcal{X}} f_j(\mathbf{x})$. It can be extracted directly from the pay-off table, but it is not robust as it can under- or over-estimate the true nadir point, see e.g., [36] for details. While for two objectives, nadir and pseudo-nadir coincide, this is not the case anymore in higher dimensions. Using this pseudo nadir point removes the need to search for extremal non-dominated solutions on the Pareto front, and aligns with finding the ideal point.

The egalitarian inspiration of the KS solution is based on the assumption that all the objectives have equal importance. However, in many practical cases, the end user may want to favor a subset of objectives which are of primary importance, while still retaining the others for optimization. One way of incorporating those *preferences* [46, 50, 43] is to discard solutions with extreme values and actually solve the following problem (constrained version of the original MOP (1)):

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \quad & \{f_1(\mathbf{x}), \dots, f_m(\mathbf{x})\} \\ \text{s.t.} \quad & f_i(\mathbf{x}) \leq c_i, \quad i \in J \subset [1, \dots, m], \end{aligned}$$

with c_i 's predefined (or interactively defined) constants. Choosing a tight value (i.e., difficult to attain) for c_i may discard a large portion of the Pareto set and favor the i -th objective over the others.

Incorporating such preferences in the KS solution can simply be done by using \mathbf{c} as the disagreement point if all objectives are constrained, or by replacing the coordinates of the nadir with the c_i values. In a game theory context, this would mean that each player would state a limit of acceptance for his objective before starting the cooperative bargaining. From geometrical considerations, one may note that \mathbf{d} (and hence \mathbf{c}) does not need to be a feasible

point. As long as \mathbf{d} is dominated by the utopia point \mathbf{u} , the KS solution remains at the intersection of the Pareto front and the (\mathbf{d}, \mathbf{u}) line. An example where preferences from a decision maker are integrated this way is provided in [8].

As shown for instance in [8], by applying a monotonic rescaling to the objectives, the corresponding KS solution can move along the Pareto front. An example of a monotonic value function applied to a raw objective corresponds to taking the logarithm of one objective for a physical or modeling reason. Since Nash equilibrium are invariant to such rescaling, they are more stable than nadir or pseudo-nadir based KS solutions, but still affected. One proposed option to alleviate this effect is to use ranks instead of objective values directly. Interestingly, it induces a dependency on the parameterization of the problem. Studying these points can include the use of a measure on the input space and the use of copulas in the objective one, [8, 9], with further work needed to quantify the effect on the Pareto front. The combination with more involved preference learning methods such as [33, 5] could also be investigated.

3 Bayesian optimization algorithms for games

Solving expensive black-box games is an emerging topic in Bayesian optimization [4, 3, 42, 8]. Similarly to the single and multi-objective cases, BO leverages the probabilistic information available from the Gaussian process models to balance exploration and exploitation in searching for the game equilibrium. Other metamodeling options could be entertained, as long as they also provide mean and variance predictions at any location, plus the ability to get realizations (posterior draws) at arbitrary locations. We present here several algorithms to solve such equilibrium or solution finding problems.

For the Nash equilibrium (Section 3.1), [4] proposed an upper confidence bound-like acquisition function based on a regret function approximating the game-theoretic regret. The case of potential games [24], when a single function can summarize the game, commonly arising in problems with shared resources, is considered by [3], with a dedicated expected improvement-like acquisition function. In [42, 8], the flexible and general step-wise uncertainty reduction is applied to estimate different classes of equilibria, including Nash and KS, which we present in Section 3.2. Although not proposed in the literature (up to our knowledge), a related family of acquisition functions based on Thompson sampling is also applicable, as we describe briefly in Section 3.3.

3.1 Fixed point approaches for the Nash equilibrium

Fixed points methods, see an example in Algorithm 8, are one way to get Nash equilibrium. In the expensive case, a naïve version is simply to replace the true objectives by their GP mean prediction: \hat{f}_i , $1 \leq i \leq m$. While computationally efficient and possibly sufficient for a coarse estimation, such an approach does not address the exploration-exploitation trade-off, and may lead to poor estimates if the models are not accurate representations of the objectives.

[3] provides an acquisition function based on the upper confidence bound to estimate Nash equilibrium for continuous games. They focus on the game theoretic regret, that is the most any player can gain from deviating from the current solution. Following small deviations for each player in turns is also the underlying principle in fixed point methods. The corresponding acquisition function aims to select as next design to evaluate one achieving the minimum of the regret of the GP surrogate. But rather than optimizing directly to get the minimum objective values, these are replaced by an approximation. That is, for a player i at \mathbf{x}_i , the minimum value is replaced by the sum of the mean objective value over \mathbf{x}_{-i} with the corresponding

Algorithm 1: Pseudo-code for the fixed-point approach [48]

```

1 Require:  $m$ : number of players,  $0 < \alpha < 1$ : relaxation factor,  $n_{\max}$ : max iterations;
2 Construct initial strategy  $\mathbf{x}^{(0)}$ ;
3 while  $n \leq n_{\max}$  do
4   Compute in parallel:  $\forall i, 1 \leq i \leq m, \mathbf{z}_i^{(k+1)} = \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \hat{f}_i(\mathbf{x}_{-i}^{(k)}, \mathbf{x}_i)$ ;
5   Update :  $\mathbf{x}^{(k+1)} = \alpha \mathbf{z}^{(k+1)} + (1 - \alpha) \mathbf{x}^{(k)}$ ;
6   if  $\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$  small enough then
7      $\lfloor$  exit;
8 Ensure: For all  $i = 1, \dots, m, \mathbf{x}_i^* = \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} J_i(\mathbf{x}_{-i}^*, \mathbf{x}_i)$ 

```

standard deviation, scaled by a factor γ . These can be computed analytically for GPs for some separable covariance functions.

In the special case of potential games [24], also related to fixed point methods, the strategy of improving as much as possible one player leads to the Nash equilibrium (and not any fixed point). This is exploited by [4] who directly considers the potential. By definition, the potential $\Phi : \mathcal{X} \rightarrow \mathbb{R}$ is a function such that $f_i(\mathbf{x}', \mathbf{x}_{-i}) - f_i(\mathbf{x}'', \mathbf{x}_{-i}) = \Phi(\mathbf{x}', \mathbf{x}_{-i}) - \Phi(\mathbf{x}'', \mathbf{x}_{-i})$, $\forall i, \mathbf{x}', \mathbf{x}''$. As the available measurements are on the f_i 's and not the potential directly, constructing the GP of the potential Φ involves taking into account integral and gradient operators. Once the GP is fitted, they proposed an acquisition function similar to expected improvement [37].

3.2 Stepwise uncertainty reduction

Step-wise uncertainty reduction (SUR) methods, see, e.g., [22, 49, 12], require an uncertainty measure Γ about a quantity of interest, here the game equilibrium. Then new evaluations are selected sequentially in order to reduce this uncertainty, leading to an accurate identification of the quantity of interest. This framework is well-adapted to games as it allows to integrate various learning tasks into a single goal (for instance, the definition of KS type solutions involves several other unknown quantities, such as the ideal and nadir points).

For any multivariate function \mathbf{f} defined over \mathcal{X} , with image \mathbb{Y} , denote by $\Psi : \mathbb{Y} \rightarrow \mathbb{R}^m$ the mapping that associates its equilibrium (Nash, KS or other). When the multivariate function is replaced by a GP emulator conditioned on n observations $(\mathbf{f}(\mathbf{x}^{(1)}), \dots, \mathbf{f}(\mathbf{x}^{(n)}))$, Ψ is defined on the corresponding distribution $\mathbf{Y}_n(\cdot)$. Then $\Psi(\mathbf{Y}_n)$ is a random vector (of unknown distribution). Loosely speaking, the spread of the corresponding distribution in the objective space reflects the uncertainty on the solution location, given by the uncertainty measure $\Gamma(\mathbf{Y}_n)$. One measure of variability of a vector is the determinant of its covariance matrix [18]: $\Gamma(\mathbf{Y}_n) = \det[\text{cov}(\Psi(\mathbf{Y}_n))]$, while other information theoretic measures could be used. For instance, this could be the conditional entropy as in [49].

The SUR strategy greedily chooses the next observation that reduces the most this uncertainty:

$$\max_{\mathbf{x} \in \mathcal{X}} \{\Gamma(\mathbf{Y}_n) - \Gamma(\mathbf{Y}_{n,\mathbf{x}})\},$$

where $\mathbf{Y}_{n,\mathbf{x}}$ is the GP conditioned on $\{\mathbf{f}(\mathbf{x}^{(1)}), \dots, \mathbf{f}(\mathbf{x}^{(n)}), \mathbf{f}(\mathbf{x})\}$. However, such an ideal strategy is not tractable as it would require evaluating all $\mathbf{f}(\mathbf{x})$ while maximizing over \mathcal{X} .

A more manageable strategy is to consider the expected uncertainty reduction: $\Gamma(\mathbf{Y}_n) - \mathbb{E}_{\mathbf{Y}_n(\mathbf{x})}[\Gamma(\mathbf{Y}_{n,\mathbf{x}})]$, where $\mathbb{E}_{\mathbf{Y}_n(\mathbf{x})}$ denotes the expectation taken over $\mathbf{Y}_n(\mathbf{x})$. Removing the

constant term $\Gamma(\mathbf{Y}_n)$, the policy is defined with

$$\mathbf{x} \in \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ J(\mathbf{x}) = \mathbb{E}_{\mathbf{Y}_n(\mathbf{x})} [\Gamma(\mathbf{Y}_{n,\mathbf{x}})] \right\}.$$

Informally speaking, the variation of Ψ is partly caused by not knowing the precise localization of the solution of interest, plus eventually from not knowing the values of the anchor points. The realizations corresponding to these specific points would either be for designs close to the solution of the GP means, or for designs with large variance. Hence $J(\mathbf{x})$ defines a trade-off between exploration and exploitation, as well as a trade-off between the different learning tasks (\mathbf{d} and \mathbf{u} points, Nash equilibrium, Pareto front). Hence the BO optimization loop is fully defined (see Chapter 11, Algorithm 1) and the difficulty boils down to evaluating J efficiently.

A crucial aspect in enabling the use of the SUR strategy is the ability to generate realizations or estimate \mathbf{Y}_n efficiently. When \mathcal{X} is discrete (or can be discretized), the approaches proposed in [42, 8] relies on the use of conditional simulations, i.e., joint posterior samples on \mathcal{X} or a well chosen subset of it, coupled with fast update formulas for the resulting ensembles, following [13]. This approach hinges on the discretization size used, which needs to remain in the thousands due to the cubic computational complexity in the number of samples. The targeted solution is then computed on each realization, resulting in samples from \mathbf{Y}_n . Note that for discrete Nash equilibrium computations, all combinations of the selected strategies for each player must be simulated at. An alternative solution, applicable to continuous \mathcal{X} would rely on approximated sample paths with a closed form expression as used, e.g., by [38]. Again, the appropriate solution can be obtained from the samples, which is inexpensive compared to using the expensive black-box.

The resulting criterion for KSE estimation automatically interleaves improving the estimation of the specific quantities involved, by balancing between estimation of the Nash equilibrium, of the objective-wise minima for the ideal point and of the intersection with the Pareto front. A different use of the posterior distribution is also possible, with Thompson sampling.

3.3 Thompson sampling

Over the past decade, Thompson sampling (TS) [47] has become a very popular algorithm to solve sequential optimization problems. In a nutshell, TS proceeds by sequentially sampling from the posterior distribution of the problem solution, allowing for efficiently addressing exploration-exploitation trade-offs. In a GP-based setting, this is simply achieved by sampling from the posterior distribution of the objectives and choosing as the next sampling point the input that realizes the equilibrium on the sample.

As it would not directly try to pinpoint the ideal point (say), it cannot be directly transposed for estimating KS solutions. Still, it may serve to estimate Nash equilibria. However, up to our knowledge this solution has not been studied yet (the closest proposition is an algorithm in [42] based on the probability of realizing the optimum). Note that the same approximation of the samples as for SUR can be used. Next we illustrate the results of this latter on a practical application.

4 Application example: engineering test case

We consider the switching ripple suppressor design problem for voltage source inversion in powered systems, which was proposed by [53], and also tested in [27]. The device is composed of three components, one with three inductors, one with parallel LC resonant branches and

a capacitor branch. Arguing that this type of problem is naturally many objective, [53] describes several types of conflicting objectives: harmonic attenuations of the switching ripple suppressor under different frequencies, power factor, resonance suppression and inductor cost.

In the proposed problem, where n_r is the number of resonant branches, there are $n_r + 4$ variables and $n_r + 1$ objectives: variables are inductors L_1, L_2, L_3 that can be related to the inductor cost objective (f_{n_r+1}), and variables C_1, \dots, C_{n_r} related to the goal of attenuating the harmonics at resonant frequencies (f_1, \dots, f_{n_r}). The last variable is C_f . Note that five additional design constraints can also be added. The cost of evaluation is here negligible, but it showcases a realistic application example where the true Pareto front is unknown.

We start by estimating the Nash-Kalai-Smorodinsky equilibrium in the $n_r = 4$ case; with the following partitioning: $(f_1 : C_1, C_f), (f_2 : C_2), (f_3 : C_3), (f_4 : C_4), (f_5 : L_1, L_2, L_3)$. For the SUR strategy, we discretize each input subspace using Latin hypercube samples of sizes (26, 11, 11, 11, 51) and use a budget of 80 evaluations for the initial design followed by 70 sequential iterations. Note that this discretization of the input space with the partitioning results in $\approx 2 \times 10^6$ possible combinations for the Nash equilibrium. To keep it tractable for the SUR procedure, filtering is performed to select 1000 solutions to actually draw posterior realizations, before evaluating the criterion on the 200 most promising candidates. More details on the procedure can be found in [42, 8]. The result obtained using the **GPGame** package [41] is described in Figures 2 and 3. It appears that the first four objectives are highly correlated, which does not affect the estimation of the solution. Most of the initial designs (67/80) are dominated, while this is the case for two thirds (47/70) of the added ones, possibly due to the search of the inefficient Nash equilibrium. These new designs are all concentrated in the same region of the output space. The NE is relatively close to the Pareto front in this case, with the obtained solution dominating it.

Then we consider the regular Kalai-Smorodinsky solution, whose results are in Figures 4 and 5. This time the input space is discretized with 10^6 possible solutions, uniformly sampled. Again, most initial design are dominated (72/80), but this is the case for only a third (28/70) of the sequentially added ones. Figure 5 shows that the KS solution is almost a straight line, an indication of the centrality of this solution. It can be seen that, in the sequential procedure, new points are added at the extremities of the Pareto front to locate the nadir, as well as in the central part.

Choosing between the obtained NKS and KS solutions now depends on the viewpoint. The spread of the latter is much broader in the objective space than for NKS. Note that the same range is used on all the Figures. To further ease comparison, we added the KS solution on Figure 2 and the Nash and NKS ones on Figure 4, but there were not part of the solutions evaluated in their corresponding BO loops, nor their targets. In case the range of objectives is relevant to the decision maker, the centrality and equity of benefits on all objectives may be appealing to the decision maker. If instead the territory splitting for the objectives is important, the relative scaling of the objectives irrelevant, then the NKS solution would be preferred (also over staying on the Nash equilibrium). As a perspective, it would be pertinent to additionally take into account the constraints defined in [53], for which the Pareto front becomes more complex.

5 What is done and what remains

We reviewed the enrichment of many-objective optimization with concepts from game theory. In the context of limited budgets of evaluations when the extensive search of the entire Pareto front is unreachable, game theory can guide the estimation of a single solution chosen for its centrality and possible additional game theoretic properties. Several BO-based method-

NKS

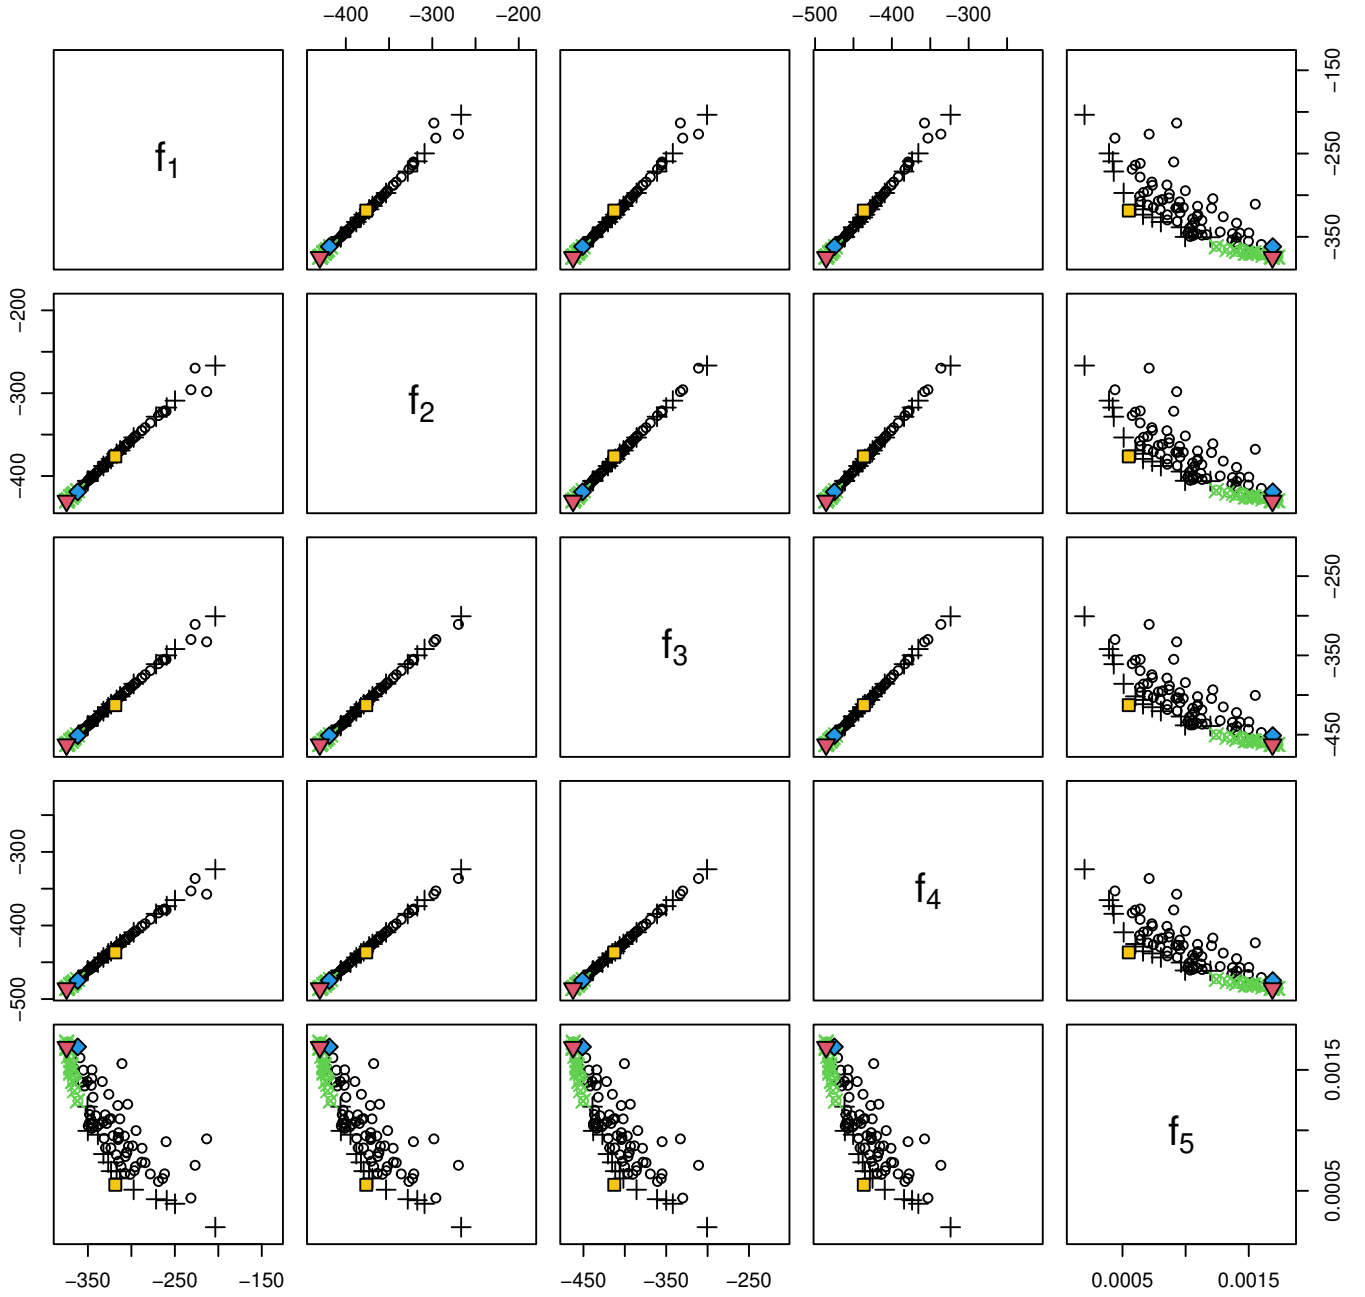


Figure 2: Scatter plot matrix of the objectives obtained when estimating the NKS solution by the SUR strategy. Black circles: dominated initial designs, black +: nondominated initial designs, green crosses: sequentially added solutions (circles if dominated), blue triangles: estimated Nash equilibrium, and red diamond: estimated NKS solution. The KS solution (yellow square) is marked for reference only.




Figure 3: Parallel coordinates plot corresponding to the objectives of Figure 2. Colors are the same, dominated solutions are in dotted lines while the Nash equilibrium is marked by the large blue dashed line and the estimated NKS solution by the large red line.

ologies have been developed for estimating the corresponding solution, which can scale to many-objective setups. As this is an emerging topic in BO, it seems that many algorithmic developments could be proposed, for instance based on Thompson sampling or on entropy search.

For the Nash and NKS solutions, a partitioning of the input space is necessary, while it is seldom included in the multi-objective context. A perspective is to randomize this partitioning when it does not come directly from the problem definition. The possible advantage would be to focus on a set of different Nash equilibria, all with stationarity properties, providing more than a single solution.

As with the engineering example we presented, there are usually additional constraints in the problem definition. In the SUR procedure described above, dealing with constraints could amount to filtering unfeasible solutions on the realizations. Other options could include the use of the probability of feasibility as in [44]. For more details on the handling of constraints in BO, we refer to [34] and references therein. Then the difference between an objective and a constraint is sometimes fuzzy. Especially when the constraints can be modified, taking them as objectives instead, akin to multiobjectivization [31], could help finding feasible solutions (or work in the case the constrained problem is unfeasible). Conversely, prioritizing some objectives [16] while defining constraints on others can help control the trade-offs when moving on the Pareto front.

References

- [1] Majid Abdolshah, Alistair Shilton, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Multi-objective Bayesian optimisation with preferences over objectives. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 12235–12245. Curran Associates, Inc., 2019.

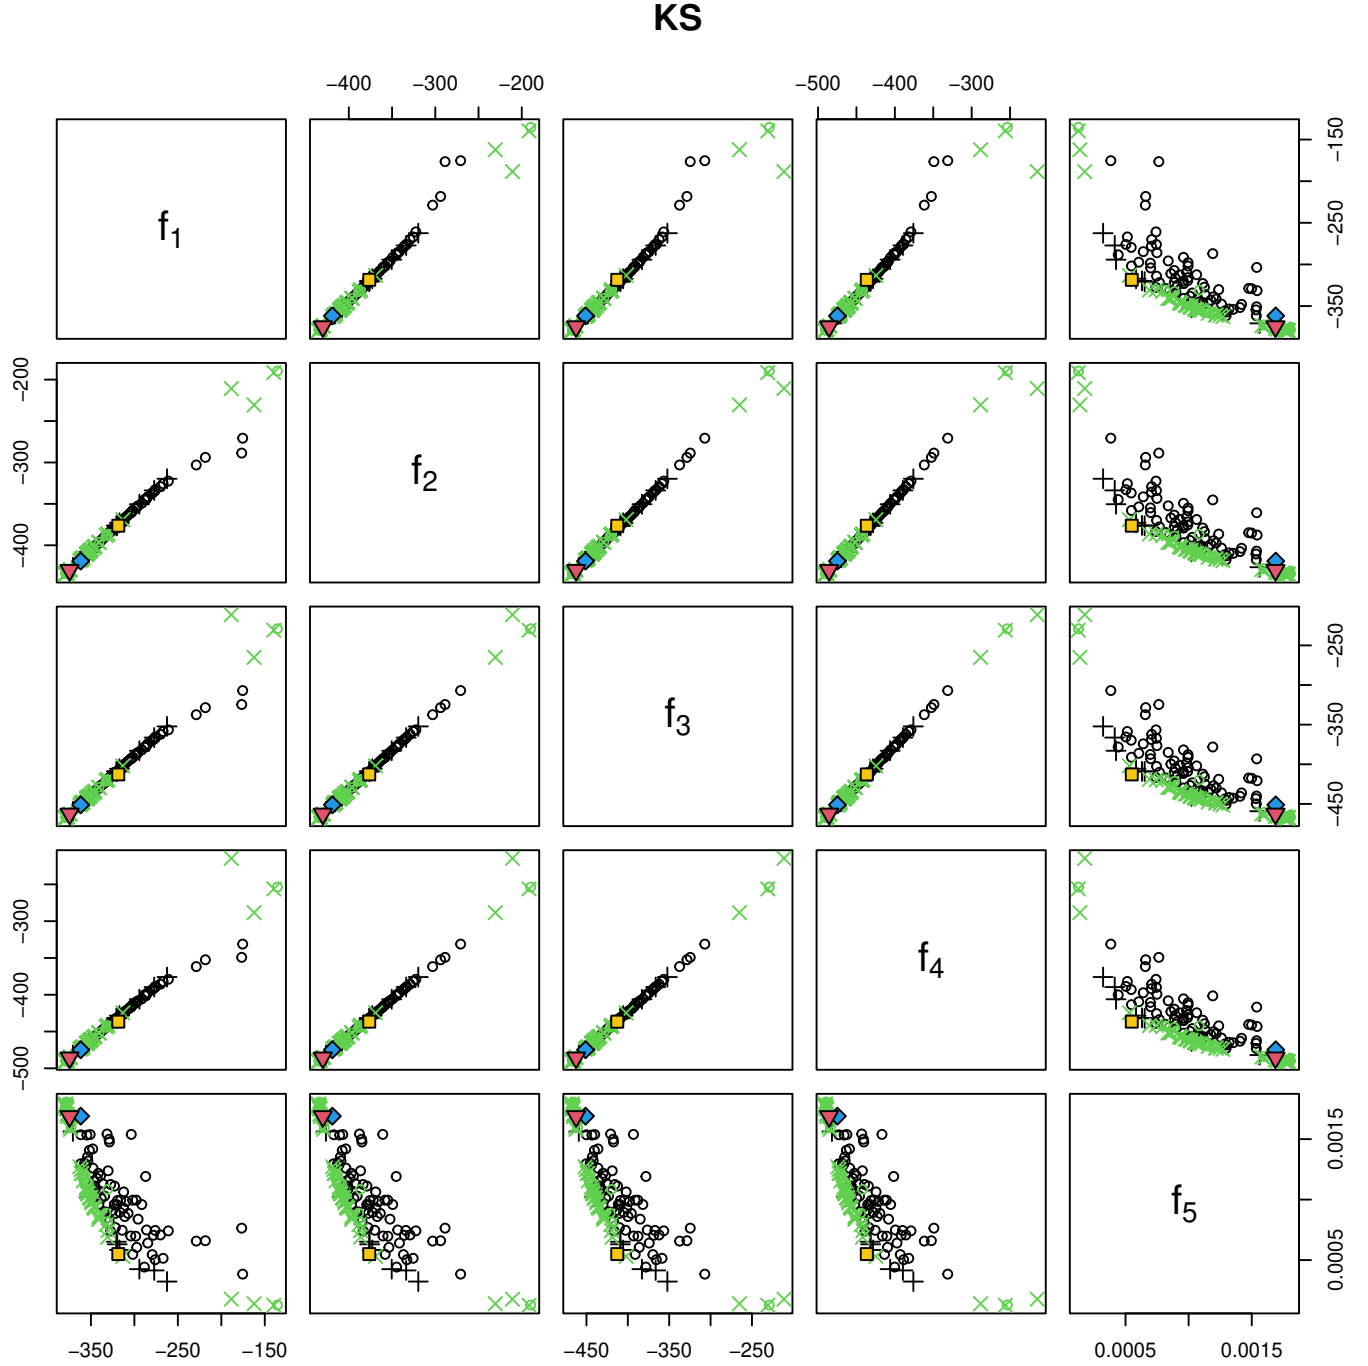


Figure 4: Scatter plot matrix of the objectives obtained when estimating the KS solution by the SUR strategy. The legend is as for Figure 2 except that the yellow square indicates the KS solution. Here the Nash solution (blue diamond) and NKS solution (red triangle) are marked for reference only.




Figure 5: Parallel coordinates plot corresponding to the objectives of Figure 4. Colors are the same, dominated solution are depicted with dotted lines while the estimated KS solution is marked by the large yellow line.

- [2] Rajae Aboulaich, Rachid Ellaia, Samira El Moumen, Abderrahmane Habbal, and Nouredine Moussaid. A new algorithm for approaching Nash equilibrium and Kalai Smoridinsky solution. working paper or preprint, December 2011.
- [3] Abdullah Al-Dujaili, Erik Hemberg, and Una-May O'Reilly. Approximating Nash equilibria for black-box games: A Bayesian optimization approach. *arXiv preprint arXiv:1804.10586*, 2018.
- [4] Anup Aprem and Stephen Roberts. A Bayesian optimization approach to compute Nash equilibrium of potential games using bandit feedback. *The Computer Journal*, 2018.
- [5] Raul Astudillo and Peter Frazier. Multi-attribute Bayesian optimization with interactive preference learning. In *International Conference on Artificial Intelligence and Statistics*, pages 4496–4507, 2020.
- [6] Tamer Başar. Relaxation techniques and asynchronous algorithms for on-line computation of noncooperative equilibria. *J. Econom. Dynam. Control*, 11(4):531–549, 1987.
- [7] Slim Bechikh, Lamjed Ben Said, and Khaled Ghedira. Estimating nadir point in multi-objective optimization using mobile reference points. In *IEEE Congress on Evolutionary Computation*, pages 1–9. IEEE, 2010.
- [8] Mickaël Binois, Victor Picheny, Patrick Taillardier, and Abderrahmane Habbal. The Kalai-Smorodinsky solution for many-objective Bayesian optimization. *Journal of Machine Learning Research*, 21(150):1–42, 2020.
- [9] Mickaël Binois, Didier Rulli re, and Olivier Roustant. On the estimation of Pareto fronts from the point of view of copula theory. *Information Sciences*, 324:270 – 285, 2015.

- [10] Irem Bozbay, Franz Dietrich, and Hans Peters. Bargaining with endogenous disagreement: The extended Kalai–Smorodinsky solution. *Games and Economic Behavior*, 74(1):407–417, 2012.
- [11] Noam Brown, Sam Ganzfried, and Tuomas Sandholm. Hierarchical abstraction, distributed equilibrium computation, and post-processing, with application to a champion no-limit Texas hold’em agent. In *AAAI Workshop: Computer Poker and Imperfect Information*, 2015.
- [12] Clément Chevalier, Julien Bect, David Ginsbourger, Emmanuel Vazquez, Victor Picheny, and Yann Richet. Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics*, 56(4):455–465, 2014.
- [13] Clément Chevalier, Xavier Emery, and David Ginsbourger. Fast update of conditional simulation ensembles. *Mathematical Geosciences*, 47(7):771–789, 2015.
- [14] John P Conley and Simon Wilkie. The bargaining problem without convexity: Extending the egalitarian and Kalai-Smorodinsky solutions. *Economics Letters*, 36(4):365–369, 1991.
- [15] Jean-Antoine Désidéri. Cooperation and competition in multidisciplinary optimization. *Computational Optimization and Applications*, 52(1):29–68, 2012.
- [16] Jean-Antoine Desideri. Platform for prioritized multi-objective optimization by metamodel-assisted Nash games. Research Report RR-9290, Inria Sophia Antipolis, September 2019.
- [17] Jean-Antoine Désidéri, Régis Duvigneau, and Abderrahmane Habbal. *Multiobjective Design Optimization Using Nash Games*, pages 583–641. 2014.
- [18] Valerii Vadimovich Fedorov. *Theory of Optimal Experiments*. Elsevier, 1972.
- [19] Jonathan E Fieldsend. Enabling dominance resistance in visualisable distance-based many-objective problems. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion*, pages 1429–1436, 2016.
- [20] David Gaudrie, Rodolphe Le Riche, Victor Picheny, Benoit Enaux, and Vincent Herbert. Targeting solutions in Bayesian multi-objective optimization: sequential and batch versions. *Annals of Mathematics and Artificial Intelligence*, 88(1):187–212, 2020.
- [21] David Gaudrie, Rodolphe Le Riche, Victor Picheny, Benoit Enaux, and Vincent Herbert. Budgeted multi-objective optimization with a focus on the central part of the Pareto front—extended version. *arXiv preprint arXiv:1809.10482*, 2018.
- [22] Donald Geman and Bruno Jedynak. An active testing model for tracking roads in satellite images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1):1–14, 1996.
- [23] R. Gibbons. *Game Theory for Applied Economists*. Princeton University Press, Princeton, NJ, 1992.
- [24] David González-Sánchez and Onésimo Hernández-Lerma. A survey of static and dynamic potential games. *Science China Mathematics*, 59(11):2075–2102, 2016.

- [25] Abderrahmane Habbal and Moez Kallel. Neumann–Dirichlet Nash strategies for the solution of elliptic Cauchy problems. *SIAM Journal on Control and Optimization*, 51(5):4066–4083, 2013.
- [26] Jussi Hakanen and Joshua D Knowles. On using decision maker preferences with ParEGO. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 282–297. Springer, 2017.
- [27] Cheng He, Ye Tian, Handing Wang, and Yaochu Jin. A repository of real-world datasets for data-driven evolutionary multiobjective optimization. *Complex & Intelligent Systems*, pages 1–9, 2019.
- [28] Jens Leth Hougaard and Mich Tvede. Nonconvex n-person bargaining: efficient maxmin solutions. *Economic Theory*, 21(1):81–95, 2003.
- [29] E. Kalai and M. Smorodinsky. Other solutions to Nash’s bargaining problem. *Econometrica*, 43:513–518, 1975.
- [30] Christian Kanzow and Daniel Steck. Augmented Lagrangian methods for the solution of generalized Nash equilibrium problems. *SIAM Journal on Optimization*, 26(4):2034–2058, 2016.
- [31] Joshua D Knowles, Richard A Watson, and David W Corne. Reducing local optima in single-objective problems by multi-objectivization. In *International conference on evolutionary multi-criterion optimization*, pages 269–283. Springer, 2001.
- [32] Marc Lanctot, Richard Gibson, Neil Burch, Martin Zinkevich, and Michael Bowling. No-regret learning in extensive-form games with imperfect recall. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1035–1042, 2012.
- [33] John R Lepird, Michael P Owen, and Mykel J Kochenderfer. Bayesian preference elicitation for multiobjective engineering design optimization. *Journal of Aerospace Information Systems*, 12(10):634–645, 2015.
- [34] Benjamin Letham, Brian Karrer, Guilherme Ottoni, Eytan Bakshy, et al. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 2018.
- [35] Shu Li and Tamer Başar. Distributed algorithms for the computation of noncooperative equilibria. *Automatica J. IFAC*, 23(4):523–533, 1987.
- [36] Kaisa Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media, 2012.
- [37] Jonas Mockus. On bayesian methods for seeking the extremum. In Guri I. Marchuk, editor, *Optimization Techniques, IFIP Technical Conference, Novosibirsk, USSR, July 1-7, 1974*, volume 27 of *Lecture Notes in Computer Science*, pages 400–404. Springer, 1974.
- [38] Mojmir Mutny and Andreas Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9005–9016. Curran Associates, Inc., 2018.

- [39] Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. A flexible framework for multi-objective Bayesian optimization using random scalarizations. In *Uncertainty in Artificial Intelligence*, pages 766–776. PMLR, 2020.
- [40] Jacques Periaux, Felipe Gonzalez, and Dong Seop Chris Lee. *Evolutionary optimization and game strategies for advanced multi-disciplinary design: applications to aeronautics and UAV design*, volume 75. Springer, 2015.
- [41] Victor Picheny and Mickael Binois. *GPGame: Solving Complex Game Problems using Gaussian Processes*, 2020. R package version 1.2.0.
- [42] Victor Picheny, Mickael Binois, and Abderrahmane Habbal. A Bayesian optimization approach to find Nash equilibria. *Journal of Global Optimization*, 73(1):171–192, 2019.
- [43] Robin C Purshouse, Kalyanmoy Deb, Maszatul M Mansor, Sanaz Mostaghim, and Rui Wang. A review of hybrid evolutionary multiple criteria decision making methods. In *2014 IEEE congress on evolutionary computation (CEC)*, pages 1147–1154. IEEE, 2014.
- [44] Matthias Schonlau, William J Welch, and Donald R Jones. Global versus local search in constrained optimization of computer models. *Lecture Notes-Monograph Series*, pages 11–25, 1998.
- [45] Joshua D. Svenson. *Computer Experiments: Multiobjective Optimization and Sensitivity Analysis*. PhD thesis, The Ohio State University, 2011.
- [46] Lothar Thiele, Kaisa Miettinen, Pekka J Korhonen, and Julian Molina. A preference-based evolutionary algorithm for multi-objective optimization. *Evolutionary computation*, 17(3):411–436, 2009.
- [47] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [48] Stanislav Uryas’ev and Reuven Y Rubinstein. On relaxation algorithms in computation of noncooperative equilibria. *IEEE Transactions on Automatic Control*, 39(6):1263–1267, 1994.
- [49] Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509–534, 2009.
- [50] Handing Wang, Markus Olhofer, and Yaochu Jin. A mini-review on preference modeling and articulation in multi-objective optimization: current status and challenges. *Complex & Intelligent Systems*, 3(4):233–245, 2017.
- [51] Andrzej P Wierzbicki. The use of reference objectives in multiobjective optimization. In *Multiple criteria decision making theory and application*, pages 468–486. Springer, 1980.
- [52] Xingyi Zhang, Ye Tian, and Yaochu Jin. A knee point-driven evolutionary algorithm for many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 19(6):761–776, 2014.
- [53] Zhixiong Zhang, Cheng He, Jie Ye, Jinbang Xu, and Linqiang Pan. Switching ripple suppressor design of the grid-connected inverters: A perspective of many-objective optimization with constraints handling. *Swarm and evolutionary computation*, 44:293–303, 2019.