



HAL
open science

Shared active subspace for multivariate vector-valued functions

Khadija Musayeva, Mickael Binois

► **To cite this version:**

Khadija Musayeva, Mickael Binois. Shared active subspace for multivariate vector-valued functions. 2023. hal-04372410

HAL Id: hal-04372410

<https://hal.science/hal-04372410>

Preprint submitted on 4 Jan 2024

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.

SHARED ACTIVE SUBSPACE FOR MULTIVARIATE VECTOR-VALUED FUNCTIONS

Khadija Musayeva, Mickaël Binois

Université Côte d'Azur, Inria, CNRS, LJAD, Sophia Antipolis, France

{khadija.musayeva, mickael.binois}@inria.fr

ABSTRACT

This paper proposes several approaches as baselines to compute a shared active subspace for multivariate vector-valued functions. The goal is to minimize the deviation between the function evaluations on the original space and those on the reconstructed one. This is done either by manipulating the gradients or the symmetric positive (semi-)definite (SPD) matrices computed from the gradients of each component function so as to get a single structure common to all component functions. These approaches can be applied to any data irrespective of the underlying distribution unlike the existing vector-valued approach that is constrained to a normal distribution. We test the effectiveness of these methods on five optimization problems. The experiments show that, in general, the SPD-level methods are superior to the gradient-level ones, and are close to the vector-valued approach in the case of a normal distribution. Interestingly, in most cases it suffices to take the sum of the SPD matrices to identify the best shared active subspace.

Keywords Dimension reduction · symmetric positive definite matrices · common principal component analysis · multi-objective optimization

1 Introduction

Many problems in machine learning, optimization, uncertainty quantification and sensitivity analysis suffer from the curse of dimensionality, where the performance and the complexity of the model worsens dramatically with the number of input variables. To alleviate this problem, one is interested in dimensionality reduction techniques. For instance, in machine learning, variable/feature selection methods Guyon and Elisseeff (2003) aim to find a subset of variables so as to improve the predictive performance of a learning algorithm, and in some algorithms, such as decision trees, the variable selection is an inherent part of the learning process. The field of sensitivity analysis mostly deals with identifying the subset of inputs parameters whose uncertainty contributes significantly to that of the model output Saltelli et al. (2008); Da Veiga et al. (2021). They are focused on the effects of the initial variables and their interactions.

However, it might be the case that the model or function of interest varies the most along directions not aligned with the coordinate axes. The widely used dimensionality reduction method of principal component analysis (PCA) (also Karhunen-Loeve method) can be used to find a linear subspace of the input/output space containing the most of its variance, but, by default, it does not take into account the input-output relationship. In ecological sciences, the redundancy analysis applies PCA to the fitted values from a linear regression model to identify a subset of input parameters contributing significantly to the variation in the response matrix Legendre et al. (2011). This is based on the

assumption of a linear relationship between the response variables and the input parameters. For regression, sufficient dimension reduction techniques aim to find a lower dimensional linear subspace such that the response and the input variables are conditionally independent given the projection of these input variables onto this subspace Cook (1996). In particular, to compute such a subspace, the work of Fukumizu and Leng (2014) is based on the gradients of the regression function. The method of active subspace of Constantine et al. (2014) also assumes that one has an access to the gradients of a function of interest, on which it performs PCA-like analysis to identify the dominant directions of variability. But it is solely applicable to scalar-valued functions. This paper is focused on the vector-valued extension of this method.

To the best of our knowledge, only two works deal with the gradient-based dimensionality reduction for vector-valued functions. The method of Zahm et al. (2020) applies only to the distributions satisfying Poincaré-type inequalities, primarily a normal distribution. To identify the dominant directions of variability, it takes into account both the function gradients and the data covariance matrix. On the other hand, Ji et al. (2018) extends the active subspace method rather trivially: it is based on a linear combination of the active subspaces computed separately for each output using the method of Constantine et al. (2014), and applies only to a normally distributed data as well. In this paper, irrespective of the data generating distribution, we aim to reduce the vector-valued problem to the scalar setting either at the level of gradients, or at the level of symmetric positive definite (SPD) matrices (i.e., covariance-like matrices), computed from these gradients. The former group of methods aims to compute a single representative vector for the component functions' derivatives with respect to each input variable, such as taking their average or projecting them to a one dimensional linear subspace. The SPD-level methods range from taking the sum of SPD matrices, to tools from the common principal component analysis (Flury (1988), Trendafilov and Gallo (2021)) to find a common structure in multiple SPD matrices. Particularly, the method of Flury and Gautschi (1986) deals with the joint diagonalization of multiple SPD matrices and that of Trendafilov (2010) with stepwise estimation of common structures in those matrices. The eigenvectors obtained from these methods then represents the important directions of variability. The experiments on one synthetic and four real-world inspired optimization problems show that the latter group of methods, in particular, the sum of SPD matrices, are very effective approaches to computing a common active subspace for vector-valued functions with respect to the root-mean-square error. In particular, for a normally distributed data their performances are very close to that of the method of Zahm et al. (2020).

The paper is organized as follows. Section 2 gives a general overview of the active subspace method of Constantine et al. (2014), and that of Zahm et al. (2020) for vector-valued functions. In Section 3, we introduce and discuss our approaches to compute a shared active subspace for vector-valued functions. Section 4 discusses the numerical experiments conducted to test their effectiveness, and the results obtained. Section 5 concludes the paper with a brief discussion of possible research directions on this topic. We provide an R package implementation of the proposed methods as well as that of Zahm et al. (2020) in the supplementary material.

2 Methodological background

In what follows, we denote by $\mathcal{X} \subseteq \mathbb{R}^d$ an input/parameter space and \mathbb{P} is a probability distribution with density ρ defined on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, where $\mathcal{B}(\mathcal{X})$ are Borel sets of \mathcal{X} . \mathbf{X} stands for a random variable with the values in \mathcal{X} and $\mathbf{D}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathcal{X}$ is an n -sample distributed according to \mathbb{P}^n . Let \mathcal{H} denote the space of square-integrable functions $\mathbf{f} = (f_1, \dots, f_C) : \mathcal{X} \rightarrow \mathbb{R}^C$ with respect to ρ , where f stands for a real-valued function, with the norm $\|\mathbf{f}\|_{\mathcal{H}}^2 = \mathbb{E}_{\mathbf{X} \sim \mathbb{P}}[\mathbf{f}(\mathbf{X})\mathbf{f}(\mathbf{X})^T]$. We denote the gradient of the k -th component of \mathbf{f} , $k \in \{1, \dots, C\}$, by $\nabla f_k(\mathbf{x}) = \left[\frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}^{(1)}} \dots \frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}^{(d)}} \right]$, where $\mathbf{x}^{(i)}$ is the i -th element of the input vector \mathbf{x} , and we denote the gradient evaluations on \mathbf{D}_n by $\nabla_{\mathbf{D}_n} f_k = [(\nabla f_k(\mathbf{x}_1))^T \dots (\nabla f_k(\mathbf{x}_n))^T]^T$. \mathbf{J} stands for the Jacobian of \mathbf{f} , i.e.,

$$\mathbf{J} = \begin{bmatrix} \frac{\partial f_1(\mathbf{x})}{\partial \mathbf{x}^{(1)}} & \cdots & \frac{\partial f_1(\mathbf{x})}{\partial \mathbf{x}^{(d)}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_C(\mathbf{x})}{\partial \mathbf{x}^{(1)}} & \cdots & \frac{\partial f_C(\mathbf{x})}{\partial \mathbf{x}^{(d)}} \end{bmatrix}.$$
 With this notation at hand, the next two sections give a brief overview of the existing active subspace methods.

2.1 Active subspace

The goal of the active subspace method of Constantine et al. (2014), on which our work is based, is to find the directions of the strongest variability of f based on the eigenvalue decomposition of the SPD matrix (which can be thought of as an uncentered covariance matrix), $\mathbf{C} = \mathbb{E}_{\mathbf{X} \sim \mathbb{P}}[(\nabla f(\mathbf{X}))^T \nabla f(\mathbf{X})]$. In the obtained decomposition, $\mathbf{C} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^T$, provided that there is a sharp decay in the eigenvalues in the form $\mathbf{\Lambda} = \text{diag}(\mathbf{\Lambda}_1, \mathbf{\Lambda}_2)$, where $\mathbf{\Lambda}_1 = [\lambda_1 \dots \lambda_l]$ and $\mathbf{\Lambda}_2 = [\lambda_{l+1} \dots \lambda_d]$, we can decompose $\mathbf{W} = [\mathbf{W}_1 \mathbf{W}_2]$. The eigenvectors \mathbf{W}_1 correspond to the largest eigenvalues $\mathbf{\Lambda}_1$ and identify the directions of the strongest average variability of f , i.e., f varies stronger on the coordinates $\mathbf{x}' = \mathbf{W}_1^T \mathbf{x}$ than on the $\mathbf{x}'' = \mathbf{W}_2^T \mathbf{x}$. Then, for a fixed \mathbf{x}' , $f(\mathbf{x}) = f(\mathbf{W}_1 \mathbf{x}' + \mathbf{W}_2 \mathbf{x}'')$ can be approximated by averaging out over all values of \mathbf{x} that map to \mathbf{x}' as $\hat{f}(\mathbf{x}') = \mathbb{E}_{\mathbf{X}'' \sim \tilde{\mathbb{P}}'_{\mathbf{x}'|\mathbf{x}'}} [f(\mathbf{W}_1 \mathbf{x}' + \mathbf{W}_2 \mathbf{X}'')]$, where $\tilde{\mathbb{P}}'$ is the joint distribution of the pair $(\mathbf{X}', \mathbf{X}'')$ with the density $\rho(\mathbf{W}_1 \mathbf{x}' + \mathbf{W}_2 \mathbf{x}'')$.

Since the computation of \mathbf{C} and \hat{f} involve that of high-dimensional integrals, in practice they are estimated based on data, or on a surrogate model, see e.g., Wycoff et al. (2021). First, \mathbf{W} is estimated from the eigendecomposition of $\hat{\mathbf{C}} = \frac{1}{n} \sum_{i=1}^n (\nabla f(\mathbf{x}_i))^T \nabla f(\mathbf{x}_i)$ computed using \mathbf{D}_n . This is, in fact, similar to performing PCA on the gradient evaluations $\nabla_{\mathbf{D}_n} f$. The obtained eigenvectors $\hat{\mathbf{W}}$ are then used to approximate f as

$$\hat{f}(\mathbf{x}') = \frac{1}{N} \sum_{i=1}^N \left[f(\hat{\mathbf{W}}_1 \mathbf{x}' + \hat{\mathbf{W}}_2 \mathbf{x}_i'') \right], \quad (1)$$

where \mathbf{x}_i'' is sampled independently from the marginal $\tilde{\mathbb{P}}'_{\mathbf{x}_i''|\mathbf{x}'}$ of $\tilde{\mathbb{P}}'$ with the density $\rho(\hat{\mathbf{W}}_1 \mathbf{x}' + \hat{\mathbf{W}}_2 \mathbf{x}'')$. Under the assumption that there is an $\epsilon > 0$ such that $\|\mathbf{W} - \hat{\mathbf{W}}\| \leq \epsilon$ in the matrix 2-norm, for this (conditional expectation) approximation of f , using the Poincaré-type inequalities the authors derive an upper bound on the mean squared error as

$$\mathbb{E}_{\mathbf{X} \sim \mathbb{P}} \left[(f(\mathbf{X}) - \hat{f}(\hat{\mathbf{W}}_1^T \mathbf{X}))^2 \right] \leq K \left(1 + \frac{1}{N} \right) \left(\epsilon \left(\sum_{i=1}^l \lambda_i \right)^{1/2} + \left(\sum_{j=l+1}^d \lambda_j \right)^{1/2} \right)^2,$$

where K is a constant depending on \mathcal{X} and \mathbb{P} . Given that the dependence on N is negligible, if ϵ is small and if $[\lambda_{l+1} \dots \lambda_d]$ decays rapidly, this bound will be small. This motivates us to consider the approximation of $f(\mathbf{x})$ as $f(\hat{\mathbf{W}}_1 \mathbf{x}') = f(\hat{\mathbf{W}}_1 \hat{\mathbf{W}}_1^T \mathbf{x})$ using only the dominant eigenvectors $\hat{\mathbf{W}}_1$ of $\hat{\mathbf{W}}$. The question is then the sample size n needed to obtain an ϵ -approximation of \mathbf{W} . To this end, Holodnak et al. (2018) provides a lower bound on the sample size n in terms of ϵ and the quantities characterizing the smoothness of f . Assuming that f is smooth enough, this dependence takes the form $\mathcal{O}(1/\epsilon^2)$.

It should be noted that Lee (2019) proposed a modification to this active subspace computation by taking into account the average of gradients $\mathbf{Z} = \mathbb{E}_{\mathbf{X} \sim \mathbb{P}}[(\nabla f(\mathbf{X}))]$ as $\mathbf{C} = \mathbb{E}_{\mathbf{X} \sim \mathbb{P}}[(\nabla f(\mathbf{X}))^T \nabla f(\mathbf{X})] + \mathbf{Z}^T \mathbf{Z}$. This is shown to be an efficient dimensionality reduction approach for quadratic functions with linear trends from the perspective of optimizing the (un)explained variance of f .

2.2 Vector-valued active subspace

Under the assumption that \mathbb{P} is a Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ (or any other distribution that satisfies Poincaré-type inequalities), Zahm et al. (2020) propose to identify the directions of variability by approximating \mathbf{f} by a ridge function

of the form $\mathbf{g} \circ \mathbf{P}_r$, where $\mathbf{P}_r \in \mathbb{R}^{d \times d}$ is a projector of rank r (ideally $r \ll d$), and $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^C$, so that $\|\mathbf{f} - \mathbf{g} \circ \mathbf{P}_r\|_{\mathcal{H}}$ is minimum. For any $\mathbf{x} \in \mathcal{X}$, such a minimizer takes the form of $\mathbb{E}_{\mathbf{X}' \sim \mathbb{P}}[\mathbf{f}(\mathbf{P}_r \mathbf{x} + (\mathbf{I} - \mathbf{P}_r) \mathbf{X}')]_{\mathcal{H}}$, i.e., the orthogonal projection of \mathbf{f} onto the set $\{\mathbf{g} \circ \mathbf{P}_r \mid \mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^C\} \cap \mathcal{H}$, where \mathbf{I} is the identity matrix of dimension d . Under the assumption that \mathbf{f} is continuously differentiable, similarly to the scalar-valued case, the upper bound on such an approximation of \mathbf{f} can be obtained thanks to the Poincaré-type inequalities as

$$\|\mathbf{f} - \mathbb{E}_{\mathbf{X}' \sim \mathbb{P}}[\mathbf{f}(\mathbf{P}_r \mathbf{x} + (\mathbf{I} - \mathbf{P}_r) \mathbf{X}')]_{\mathcal{H}}\|_{\mathcal{H}}^2 \leq \text{trace}(\Sigma(\mathbf{I} - \mathbf{P}_r)^T \mathbf{H}(\mathbf{I} - \mathbf{P}_r)), \quad (2)$$

where Σ is a (non-singular) covariance matrix and \mathbf{H} can be expressed as $\mathbb{E}_{\mathbf{X} \sim \mathbb{P}}[\mathbf{J}^T(\mathbf{X})\mathbf{J}(\mathbf{X})]$. The projector \mathbf{P}_r minimizing (2) is then derived in terms of the generalized eigenvectors $\mathbf{v}_i \in \mathbb{R}^d$ of the pair $(\mathbf{H}, \Sigma^{-1})$ as $\mathbf{P}_r = (\sum_{i=1}^r \mathbf{v}_i \mathbf{v}_i^T) \Sigma^{-1}$. Thus, the bound (2) is characterized not only by the gradients of \mathbf{f} , but also by the covariance structure of $\mathbf{X} \sim \mathbb{P}$. Similarly to the single-objective case, since 2 involves high-dimensional integrals, \mathbf{f} is approximated based on data as follows. First, \mathbf{H} is estimated as $\hat{\mathbf{H}} = \frac{1}{n} \sum_{i=1}^n \mathbf{J}^T(\mathbf{x}_i) \mathbf{J}(\mathbf{x}_i)$ based on \mathbf{D}_n . The projection $\hat{\mathbf{P}}_r$ is then computed based on the generalized eigendecomposition of the pair $(\hat{\mathbf{H}}, \Sigma^{-1})$. Finally, for any \mathbf{x} , $\mathbb{E}_{\mathbf{X}' \sim \mathbb{P}}[\mathbf{f}(\mathbf{P}_r \mathbf{x} + (\mathbf{I} - \mathbf{P}_r) \mathbf{X}')]_{\mathcal{H}}$ is estimated based on N -sample $(\mathbf{x}'_1, \dots, \mathbf{x}'_N)$ distributed according to \mathbb{P}^N as

$$\hat{\mathbf{f}}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \mathbf{f}(\hat{\mathbf{P}}_r \mathbf{x} + (\mathbf{I} - \hat{\mathbf{P}}_r) \mathbf{x}'_i). \quad (3)$$

One can notice a similar form to the single-objective case approximation (1), where now the second part being an orthogonal projection does not depend on the first part, thus the same N -sample is used for any \mathbf{x} .

Another active subspace approach for vector-valued functions is that of Ji et al. (2018). Under the assumption that \mathbb{P} is the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, their goal is to estimate the marginal distributions of $f_k(\mathbf{X})$, $1 \leq k \leq C$, based on the active subspace approach. This work consists of, first, computing the leading eigenvectors \mathbf{W}_k for each f_k separately based on the work of Constantine et al. (2014) as discussed in Section 2.1, then finding the matrix $\mathbf{V} \in \mathbb{R}^{d \times r}$, where r is the dimension of the active subspace assumed to be the same for all f_k , so that $\mathbf{W}_k^T \mathbf{V} = \mathbf{I}$, i.e., \mathbf{V} is common to all component functions. Such a matrix gives the identity $f_k(\mathbf{x}) = f_k(\mathbf{V} \mathbf{W}_k^T \mathbf{x})$. Then, under the assumption that $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{W}_k^T \mathbf{X}$ also follows the standard Gaussian distribution and thus $f_k(\mathbf{X})$ is distributed in the same way as $f_k(\mathbf{V} \mathbf{W}_k^T \mathbf{X})$.

In the next section, we propose new approaches to shared active subspace computation for vector-valued functions which do not depend on the data generating distribution.

3 New baselines to vector-valued active subspace

The straightforward vector-valued extension of Constantine et al. (2014) consists of treating each component of $\mathbf{f} = (f_1, \dots, f_C)$ as a separate problem, and computing an active subspace for each separately. But this is limiting (unless some functions share exactly the same active directions, which is a rather restrictive assumption), since any joint analysis would need to be performed on the concatenation of all active subspaces.

Our goal here is to find eigenvectors $\hat{\mathbf{W}}_j$, $1 \leq j \leq d$, common to all objectives f_1, \dots, f_C , irrespective of the distribution \mathbb{P} , so as to minimize

$$\sum_{k=1}^C \left(\frac{1}{n} \sum_{i=1}^n (f_k(\mathbf{x}_i) - f_k(\hat{\mathbf{W}}_j \hat{\mathbf{W}}_j^T \mathbf{x}_i))^2 \right)^{1/2}. \quad (4)$$

We do this by transforming the vector-valued problem into the scalar-valued one of Section 2.1 either at the level of the gradients, where the Jacobian at each point of \mathbf{D}_n is replaced by a d -dimensional vector, or at the level of the SPD

matrices by manipulating the component-wise SPD matrices $\hat{\mathbf{C}}_k = \frac{1}{n} \sum_{i=1}^n (\nabla f_k(\mathbf{x}_i))^T \nabla f_k(\mathbf{x}_i)$, $1 \leq k \leq C$. In a sense, these approaches are in the spirit of the work of Ji et al. (2018), where the combination is now performed at the level of gradients ∇f_k or SPD matrices $\hat{\mathbf{C}}_k$, $1 \leq k \leq C$, instead of eigenspaces.

3.1 Manipulation at the level of gradients

In this section, we propose several approaches to reduce the Jacobian evaluated at each point \mathbf{x} in \mathbf{D}_n to a d -dimensional vector $\mathbf{u}(\mathbf{x})$. Once such a vector is computed, one can proceed as in Section 2.1.

3.1.1 Linear projection of gradients

One of the ways of doing so is to project the derivatives of each objective with respect to (wrt) each input parameter $\mathbf{x}^{(j)}$, $j \in \{1, \dots, d\}$, to the one-dimensional subspace that captures the most of the variance in these derivatives, i.e., to perform PCA-like analysis. In detail, we combine the derivative evaluations on \mathbf{D}_n of each component function

f_k , $1 \leq k \leq C$, wrt the j -th parameter into a single matrix $\mathbf{J}_j = \begin{bmatrix} \frac{\partial f_1(\mathbf{x}_1)}{\partial \mathbf{x}_1^{(j)}} & \cdots & \frac{\partial f_C(\mathbf{x}_1)}{\partial \mathbf{x}_1^{(j)}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1(\mathbf{x}_n)}{\partial \mathbf{x}_n^{(j)}} & \cdots & \frac{\partial f_C(\mathbf{x}_n)}{\partial \mathbf{x}_n^{(j)}} \end{bmatrix}$. Then, for each \mathbf{J}_j , we

find a direction $\mathbf{w}_j \in \mathbb{R}^{C \times 1}$ capturing the most of the variability in \mathbf{J}_j from the eigendecomposition of the SPD matrix $\frac{\mathbf{J}_j^T \mathbf{J}_j}{n} \in \mathbb{R}^{C \times C}$ and approximate \mathbf{J}_j as $\mathbf{J}_j \mathbf{w}_j$. Then, the i -th row in the matrix $\mathbf{A} = [\mathbf{J}_1 \mathbf{w}_1 \dots \mathbf{J}_d \mathbf{w}_d] \in \mathbb{R}^{n \times d}$ combining these rank-1 projections can be seen as a representative vector for the Jacobian evaluated at \mathbf{x}_i . The eigendecomposition of $\mathbf{A}^T \mathbf{A} / n$ yields the eigenvectors common to all objectives. For a given input variable, this approach will give the most of the weight to the component function that highly varies wrt this variable.

3.1.2 Convex hull of gradients

We notice that the problem we are interested in bears similarity to finding common descent direction for multiple objectives in the multi-objective optimization task. Désidéri (2012) shows that if there is no zero convex combination of gradient vectors $\nabla f_k(\mathbf{x})$, i.e., if there does not exist a positive $\alpha \in \mathbb{R}_+^C$ with $\alpha \alpha^T = 1$ satisfying $\alpha \mathbf{J} = \mathbf{0}$, then a descent direction from this point common to all gradient vectors exists. It is calculated based on the vector $\mathbf{u}(\mathbf{x}) \in \mathbb{R}^d$ satisfying $\mathbf{J} \mathbf{u}(\mathbf{x})^T \geq \mathbf{0}$. Furthermore, such a vector can be found by computing the minimum (Euclidean) norm element of the convex hull of the set $\{\nabla f_1(\mathbf{x}), \dots, \nabla f_C(\mathbf{x})\}$,

$$\mathbf{U} = \left\{ \mathbf{u}(\mathbf{x}) \in \mathbb{R}^d \mid \mathbf{u}(\mathbf{x}) = \sum_{k=1}^C \alpha_k \nabla f_k(\mathbf{x}), \forall k, \alpha_k \geq 0, \sum_{k=1}^C \alpha_k = 1 \right\}.$$

Here, we are not interested in the existence of the common descent direction since we are not performing optimization, but only in finding the vector $\mathbf{u}(\mathbf{x})$ at each point \mathbf{x} . A more egalitarian approach would be to simply consider the average vector

$$\mathbf{u}(\mathbf{x}) = \frac{1}{C} \sum_{k=1}^C \nabla f_k(\mathbf{x}),$$

which is also an element of \mathbf{U} . However, it is not necessarily positively correlated with the gradient vectors, as it is the case for the minimum norm element of \mathbf{U} . On the other hand, if $\mathbf{u}(\mathbf{x})$ is the average of the normalized gradients, $\mathbf{J}'_k = \nabla f_k(\mathbf{x}) / \|\nabla f_k(\mathbf{x})\|_2$, it will satisfy $\mathbf{J}' \mathbf{u}(\mathbf{x})^T \geq \mathbf{0}$, where \mathbf{J}'_k stands for the k -th row of \mathbf{J}' . However, normalizing the gradients does not suit the context of this work, because it changes their magnitudes which determine the active subspace. Given an input variable, if the linear projection of the gradients favors the highly varying objective and the minimum norm element of the convex hull the least varying one, the average will provide a compromise solution between these two.

In this approach, a shared eigenspace is computed from the eigendecomposition of $\frac{1}{n} \sum_{i=1}^n \mathbf{u}(\mathbf{x}_i) \mathbf{u}(\mathbf{x}_i)^T$.

3.2 Manipulation at the level SPD matrices

We can compute a common eigenspace by combining the component matrices $\hat{\mathbf{C}}_k$, $1 \leq k \leq C$, into a single SPD matrix $\hat{\mathbf{C}}$ and eigendecomposing it. A more involved approach is related to the common principal component analysis, a generalization of PCA to several covariance matrices where the goal is to find common eigenvectors to all of them.

3.2.1 Sum of SPD matrices

The straightforward approach to obtain a single common SPD matrices is to consider the sum $\hat{\mathbf{C}} = \sum_{k=1}^C \hat{\mathbf{C}}_k$, which is itself an SPD matrix. Notice that, for the standard normal distribution, this method will provide similar results to that in (3), since in this case, the projector becomes $P_r = \sum_{i=1}^r \mathbf{v}_i \mathbf{v}_i^T$ which is the same as the sum just mentioned. In fact, Krzanowski (1979) has communicated the effectiveness of the average of the sample covariance matrices in the estimation of approximate common principal components. It should be noted that any scaling $\alpha \hat{\mathbf{C}}$, with $\alpha > 0$, will only scale the eigenvalues as Λ/α leaving intact the eigenvectors, consequently, it has no effect on the computation of the dominating eigenspace.

3.2.2 Joint diagonalization of SPD matrices

A more involved approach to find a common structure in $\{\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_C\}$ is based on the joint diagonalization of these matrices Flury and Gautschi (1986); Flury (1987). This method has been used for estimating the common principal components for multiple groups and a series of works has been dedicated to this line of research (Beaghen (1997); Trendafilov (2010), Browne and McNicholas (2014), Trendafilov and Gallo (2021) and references therein). If all $\hat{\mathbf{C}}_k$ are simultaneously diagonalizable, the equality $\mathbf{W}^T \hat{\mathbf{C}}_k \mathbf{W} = \Lambda_k$, where Λ_k is diagonal, for all k holds. But usually this is not the case and Flury and Gautschi (1986) (FG) proposed an iterative algorithm to compute $\hat{\mathbf{W}}$ such that the matrices $\hat{\mathbf{W}}^T \hat{\mathbf{C}}_k \hat{\mathbf{W}}$ are as diagonal as possible by minimizing the deviation from the diagonality,

$$\prod_{k=1}^C \left(\det(\text{diag}(\hat{\mathbf{W}}^T \hat{\mathbf{C}}_k \hat{\mathbf{W}})) / \det(\hat{\mathbf{W}}^T \hat{\mathbf{C}}_k \hat{\mathbf{W}}) \right)^{n_k}, \quad (5)$$

where $\det(\cdot)$ stands for the determinant and n_k is the number of samples in the group k , which in our case is equal to n . The quantity (5) appears as a test statistic of the simultaneous diagonalization of the covariance matrices of multiple normally distributed random variables Flury (1984).

The specifics of the FG algorithm is that although it finds a common structure in $\hat{\mathbf{C}}_k$'s, it does not produce any canonical ordering on $\hat{\mathbf{W}}$. In other words, although the eigenvectors are common to all components, the decay structure is not necessarily the same for all Λ_k . Hence, to identify the active subspace, we need to induce an ordering on the columns of $\hat{\mathbf{W}}$. Although different weighting approaches can be applied to induce an ordering on the eigenvalues from the FG algorithm, to keep it simple, in this paper we identify the dominating subspace based on the ordering induced by the sum $\sum_{k=1}^C \Lambda_k$.

3.2.3 Stepwise estimation of eigenvectors

In fact, due to the arbitrariness of the ordering of eigenvectors, the FG method has been considered unsuitable for the dimensionality reduction task. Instead, Trendafilov (2010) proposed an approach to find approximate common eigenvectors in a stepwise manner so that all groups share approximately the same decay structure. He shows that the

minimization problem (5) is equivalent to

$$\sum_{k=1}^C n_k \log(\det(\text{diag}(\hat{\mathbf{W}}^T \hat{\mathbf{C}}_k \hat{\mathbf{W}}))),$$

which can be expressed in the vectorized form as

$$\sum_{k=1}^C n_k \sum_{j=1}^d \log(\hat{\mathbf{w}}_j^T \hat{\mathbf{C}}_k \hat{\mathbf{w}}_j).$$

The latter quantity can then be minimized by sequentially minimizing the d identical problems of the form

$$\sum_{k=1}^C n_k \log(\hat{\mathbf{w}}_j^T \hat{\mathbf{C}}_k \hat{\mathbf{w}}_j), \quad 1 \leq j \leq d.$$

More precisely, each $\hat{\mathbf{w}}_j$ is found in the orthogonal subspace of the previous solution $\hat{\mathbf{w}}_{j-1}$, thus yielding an ordering on the decreasing importance of these vectors, and this ordering holds simultaneously for all $\{\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_C\}$. The approach of Trendafilov yields the same result as that of FG method, if the eigenvalues Λ_k decreases simultaneously for all k groups.

4 Numerical Experiments

Our experiments include one synthetic, and the following four real-world inspired optimization problems: automobile car side-impact problem Deb et al. (2009), conceptual marine design problem Parsons and Scott (2004); Tanabe and Ishibuchi (2020), penicillin production Liang and Lai (2021) and switching ripple suppressor design problem Zhang et al. (2019); He et al. (2020). In the following sections, we detail the experimental setup and give a general overview of the problems considered.

4.1 Setup

We performed two sets of experiments with different data distributions. In the first set, to compare the results with those of the method of Zahm et al. (2020), the data were sampled from the standard normal distribution. In the second set, since the input variables in the test problems are independent and all values are equally worth, we use the uniform distribution on $[-1, 1]^d$. In each problem, we have $\mathcal{X} = [u_1, l_1] \times \dots \times [u_d, l_d]$, and this requires an additional step to map the data to the hypercube $[0, 1]^d$ and apply the transformation $x^{(i)} \cdot (u_i - l_i) + l_i$ for all i . To perform this mapping, we used the sigmoid function for the normal distribution, and $(x^{(i)} + 1)/2$ for all i for the uniform one. We also mean-variance normalized the function outputs.

Active subspace is based on the gradient computations which can be a costly task with a large number of input variables. Thus to have a good estimation of the eigenspace of dimension d' , Constantine and Gleich (2015); Constantine (2015) propose to use $\alpha d' \log(d)$ number of samples, where α varies between 2 and 10. Since in our case for all problems $d < 10$, using samples of size 1000 we could compute the gradients in a matter of seconds using an ordinary laptop. We repeated the experiments 10 times. In the method of Zahm et al. (2020), for all problems, the number N of samples in Equation (3) was set to 20 (the maximum value used in the original paper).

We provide the implementation of the proposed approaches as an R package SharedAS. We compute the gradients using the R package numDeriv Gilbert et al. (2009). For the FG algorithm and the stepwise estimation of the eigenvectors we use the R package cpc Pepler (2014). The FG method, by default, returns the eigenvectors in the order of the increasing sum of the component eigenvalues. It can be estimated either by maximum likelihood (the original approach) or least squares proposed by Beaghen (1997) which we actually used. We computed the minimum element of the convex hull based on the algorithm of Sekitani and Yamamoto (1993).

Problem	number of input variables	number of objectives
Synthetic problem	3	2
Automobile car side-impact	7	11
Conceptual marine design	6	3
Penicillin production	7	5
Switching ripple suppressor design	8	5 and 10

Table 1: Summary of the test problems.

4.2 Test problems

The problems we considered are listed in Table 1.

Synthetic problem As a synthetic problem we consider a modification of the multi-objective test problem defined by Parr (2013): let $b_1 = 15x_1 - 5$, $b_2 = 15x_2$ and $b_3 = x_3$,

$$f_1(\mathbf{x}) = (b_2 - 5.1(b_1/(2\pi))^2 + 5/\pi b_1 - 6)^2 + 10((1 - 1/(8\pi)) \cos(b_1) + 1) + \sin(\pi b_3)$$

and

$$f_2(\mathbf{x}) = -\sqrt{(10.5 - b_1)(b_1 + 5.5)(b_2 + 0.5)} - \frac{1}{30} \left(b_2 - \frac{5.1}{4\pi^2} b_1^2 - 6 \right)^2 - \frac{1}{3} \left[\left(1 - \frac{1}{8\pi} \right) \cos(b_1) + 1 \right] - \cos(2\pi b_3).$$

Each of these 3-dimensional functions is rotated with a different rotation matrix: the first one is $\begin{pmatrix} -0.71 & 0.34 & 0.62 \\ 0.28 & 0.94 & -0.2 \\ -0.65 & 0.032 & -0.76 \end{pmatrix}$, and the second one $\begin{pmatrix} -0.32 & 0.84 & -0.44 \\ 0.76 & -0.058 & -0.65 \\ 0.57 & 0.54 & 0.62 \end{pmatrix}$.

Automobile car side-impact problem He et al. (2020) This problem of vehicle performance optimization involves 7 decision variables for thicknesses of car body parts (4 additional noise variables are kept fixed) and 11 objectives for the car mass, deceleration and structural integrity.

Conceptual marine design problem Tanabe and Ishibuchi (2020) This is a simplified commercial ship sizing problem with 6 input parameters, and 3 objectives which are the transportation cost, light ship weight and annual cargo capacity.

Penicillin production problem Liang and Lai (2021) is based on 7 input variables, which include culture medium volume V , biomass concentration X , glucose substrate concentration S , and three objectives which are penicillin concentration P , carbon dioxide concentration CO_2 and the reaction time t . Since the latter is a discrete quantity, we exclude it from our computation. In our experiments, everything else being equal to that of the original setting, we set the reaction time $t = 100$, and remove all the outputs where the loop over t terminates due to the extreme values of S and V . Since V , X and S are differentiable quantities we used them as outputs, along with P and CO_2 , thus dealing with a problem with 5 objectives. For this problem, we used the implementation provided on github.com/HarryQL/TuRBO-Penicillin.

Switching ripple suppressor design Zhang et al. (2019) The problem of the switching ripple suppressor design for voltage source inversion in powered system is represented by 8 input variables and 5 objectives. The input variables are the design parameters of the electronic components, and the objectives are the total cost of the inductors and the harmonics attenuations at two different resonant frequencies. As it is the case with the penicillin production problem,

here we also used some of the constraints as outputs, thus increasing the number of the objectives to 10. Overall, for this problem, we experimented separately with 5 and 10 objectives.

4.3 Results

In what follows, AG stands for the average of gradients, MCH for the minimum norm element of the convex hull of the gradients, LP for the linear projection of gradients, SSPD for the sum of the SPD matrices, SEE for the stepwise estimation of eigenvectors, FG for Flury-Gautschi algorithm, and finally, Zahm for the method of Zahm et al. (2020). The results are illustrated as boxplots in Figure 1, for the standard normal, and in Figure 2, for the uniform distribution, representing the sum of RMSEs (4) against the dimensionality j of the dominating eigenspace $\hat{\mathbf{W}}_j$ defining a shared active subspace (different for each method). Note that, the vectors in $\hat{\mathbf{W}}_j$ are ordered by their importance. Because of space constraints, for all real-world problems, we have shown only the last 4 dimensions.

In all figures, one can easily observe that, as the number of dimensions increases, the error (4) of all methods decreases, converging to zero for $\hat{\mathbf{W}}_d$, which contains all the eigenvectors, demonstrating the soundness of these methods. In the case of the normal distribution, for most problems and the dimensionality j , the method of Zahm performs the best, followed by the SPD-level methods. Both for the standard normal and the uniform distribution, the SPD-level methods outperform the gradient-level ones and perform very similarly to each other, exhibiting a greater stability. From the gradient-level methods, the minimum-norm element of the convex hull exhibited the worst performance: although being positively correlated with all the gradients, having the minimum norm, it downplays the variability. Thus, we excluded this method from the graphs of all real-world optimization problems. It is noteworthy that the simple sum of SPD matrices provides similar to or better results than SSPD and FG, suggesting that it is a very useful instructive baseline.

Sufficient summary plot (see Constantine (2015)) is a useful tool to visualize the quality of the computed active subspace. Figure 3 stands for such a plot for the synthetic problem where only a subset of methods has been shown; the rest is given in Figure 4 in Appendix A. The first two plots in all panels illustrate the outputs of each function, respectively, on the given data $\mathbf{D}_n \in \mathbb{R}^{n \times d}$, marked as black ones, and on the rotated data $\mathbf{D}_n \hat{\mathbf{W}}_1 \hat{\mathbf{W}}_1^T$, marked as red points, against the active subspace of dimension 1, $\mathbf{D}_n \hat{\mathbf{W}}_1$. The third one reflects simultaneously the pairwise plot of the original outputs of the two component functions, and their outputs on $\mathbf{D}_n \hat{\mathbf{W}}_1 \hat{\mathbf{W}}_1^T$ (red points), and $\mathbf{D}_n \hat{\mathbf{W}}_2 \hat{\mathbf{W}}_2^T$ (green points). For this synthetic problem, the second objective f_2 demonstrates greater variability than f_1 wrt all input variables: this is captured by all methods, except AG and MCH, and is reflected in the form of a univariate trend in their corresponding (the middle) scatter plots. As the dimensionality j is increased to 2, the scatter plots in the rightmost graph for all methods, except for AG and MCH, represent well the original output space. This can also be verified from the summary given in Table 2a, and is inline with the boxplots of Figure 2.

Finally, as a remark, Table 2 in the main body, and Table 3 in Appendix B, report the impact of the normalization of the function outputs when computing the gradients, on the RMSE performances of the methods: the performance improvement it provides, both in terms of the error and the variance, depends on the nature of the problem and the method considered, and the dimensionality of the active subspace. In particular, such a manipulation affects negatively the gradient-level methods (they are sensitive to the magnitudes of the gradients), whereas the SPD-level ones mostly benefit from it (they seek a common structure in the SPD matrices).

5 Conclusion and perspectives

We proposed several baseline gradient-based methods to compute a shared active subspace for multivariate vector-valued functions. These methods aim to reduce the vector-valued problem to the scalar-valued setting for which there exists an effective active subspace method, either at the level of the gradients or at the level of the SPD matrices computed from these gradients. They can be applied to any data regardless of the underlying distribution unlike the existing vector-valued approach which is based on the distributions satisfying the Poincaré-type inequalities. The experiments




Figure 1: Comparison of the proposed approaches and the method of Zahm based on data sampled from the standard normal distribution. The experiments were repeated 10 times with the sample size $n = 1000$ in each of them.




Figure 2: Comparison of the proposed approaches based on data sampled from the uniform distribution. The experiments were repeated 10 times with the sample size $n = 1000$ in each of them.

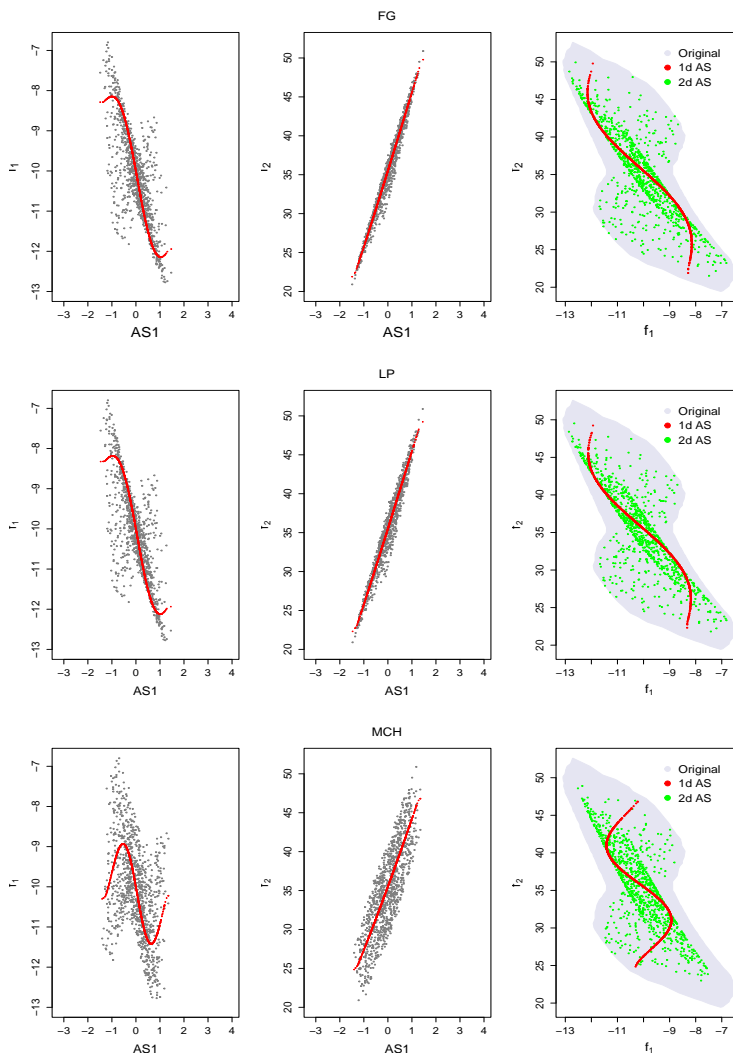


Figure 3: Sufficient summary plots of a subset of methods for the synthetic problem applied to uniformly distributed data. AS1 stands for the active subspace of dimension 1, and f_1, f_2 are the two objectives.

show that regardless the dimensionality of the active subspace, the SPD-level methods perform better, both in terms of the error and the stability, than the gradient-level ones, and close to the vector-valued approach in the case of a normal distribution. The noteworthy fact is that, in most cases, the best shared active subspace can be computed by the simple sum of the SPD matrices.

In this work we did not give any preference to objective values, neither did we take into account any particularity or structure in the problem of interest. However, in some applications one might favor certain regions in the image space of the function of interest, for which it is pertinent to consider a dedicated shared active subspace. One of the approaches to this problem may be based on the use of data depth tools which measure the centrality of a given point in a data cloud (more in Appendix C). Another perspective, inspired from Spagnol et al. (2019), would be to combine seamlessly multi-objective optimization and dimensionality reduction.

Method	j=1			j=2		
	f_1	f_2	sum	f_1	f_2	sum
AG	0.969	0.994	1.963	0.609	0.156	0.765
MCH	0.817	0.606	1.423	0.545	0.296	0.841
LP	0.739	0.210	0.949	0.129	0.182	0.311
SSPD	0.756	0.281	1.037	0.171	0.146	0.317
SEE	0.758	0.259	1.017	0.245	0.083	0.328
FG	0.746	0.188	0.934	0.166	0.155	0.321

(a)

Method	j=1			j=2		
	f_1	f_2	sum	f_1	f_2	sum
AG	0.774	0.135	0.909	0.373	0.102	0.475
MCH	1.207	1.228	2.435	1.195	1.080	2.275
LP	0.699	0.056	0.755	0.532	0.046	0.578
SSPD	0.707	0.057	0.764	0.292	0.048	0.340
SEE	0.741	0.256	0.997	0.246	0.081	0.327
FG	0.706	0.056	0.762	0.317	0.050	0.367

(b)

Table 2: Synthetic problem: the objective-wise RMSEs, f_1 , f_2 , and the sum. In (b) the gradients are computed from the original outputs, whereas in (a) from the mean-variance normalized outputs. The best performing method in (each table separately) is shown in bold.

References

- Beaghen, M. (1997). *Canonical variate analysis and related methods with longitudinal data*. Ph. D. thesis, Virginia Polytechnic Institute and State University.
- Browne, R. P. and P. D. McNicholas (2014). Estimating common principal components in high dimensions. *Advances in Data Analysis and Classification* 8, 217–226.
- Constantine, P. and D. Gleich (2015). Computing active subspaces with monte carlo.
- Constantine, P. G. (2015). *Active subspaces: Emerging ideas for dimension reduction in parameter studies*. SIAM.
- Constantine, P. G., E. Dow, and Q. Wang (2014). Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM Journal on Scientific Computing* 36(4), A1500–A1524.
- Cook, R. D. (1996). Graphics for regressions with a binary response. *Journal of the American Statistical Association* 91(435), 983–992.
- Da Veiga, S., F. Gamboa, B. Iooss, and C. Prieur (2021). *Basics and trends in sensitivity analysis: Theory and practice in R*. SIAM.
- Deb, K., S. Gupta, D. Daum, J. Branke, A. K. Mall, and D. Padmanabhan (2009). Reliability-based optimization using evolutionary algorithms. *IEEE transactions on evolutionary computation* 13(5), 1054–1074.
- Désidéri, J.-A. (2012). Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathématique* 350(5-6), 313–318.
- Flury, B. (1988). *Common principal components & related multivariate models*. John Wiley & Sons, Inc.
- Flury, B. K. (1987). Two generalizations of the common principal component model. *Biometrika* 74(1), 59–69.
- Flury, B. N. (1984). Common principal components in k groups. *Journal of the American Statistical Association* 79(388), 892–898.
- Flury, B. N. and W. Gautschi (1986). An algorithm for simultaneous orthogonal transformation of several positive definite symmetric matrices to nearly diagonal form. *SIAM Journal on Scientific and Statistical Computing* 7(1), 169–184.
- Fukumizu, K. and C. Leng (2014). Gradient-based kernel dimension reduction for regression. *Journal of the American Statistical Association* 109(505), 359–370.
- Gilbert, P., R. Varadhan, and M. P. Gilbert (2009). Package ‘numderiv’. *differential equations* 3, 203–267.

- Guyon, I. and A. Elisseeff (2003). An introduction to variable and feature selection. *Journal of machine learning research* 3(Mar), 1157–1182.
- He, C., Y. Tian, H. Wang, and Y. Jin (2020). A repository of real-world datasets for data-driven evolutionary multiobjective optimization. *Complex & Intelligent Systems* 6, 189–197.
- Holodnak, J. T., I. C. Ipsen, and R. C. Smith (2018). A probabilistic subspace bound with application to active subspaces. *SIAM Journal on Matrix Analysis and Applications* 39(3), 1208–1220.
- Ji, W., J. Wang, O. Zahm, Y. M. Marzouk, B. Yang, Z. Ren, and C. K. Law (2018). Shared low-dimensional subspaces for propagating kinetic uncertainty to multiple outputs. *Combustion and Flame* 190, 146–157.
- Krzanowski, W. (1979). Between-groups comparison of principal components. *Journal of the american statistical association* 74(367), 703–707.
- Lee, M. R. (2019). Modified active subspaces using the average of gradients. *SIAM/ASA Journal on Uncertainty Quantification* 7(1), 53–66.
- Legendre, P., J. Oksanen, and C. J. ter Braak (2011). Testing the significance of canonical axes in redundancy analysis. *Methods in Ecology and Evolution* 2(3), 269–277.
- Liang, Q. and L. Lai (2021). Scalable bayesian optimization accelerates process optimization of penicillin production. In *NeurIPS 2021 AI for Science Workshop*.
- Mosler, K. (2013). *Depth Statistics*, pp. 17–34. Springer Berlin Heidelberg.
- Parr, J. (2013). *Improvement criteria for constraint handling and multiobjective optimization*. Ph. D. thesis, University of Southampton.
- Parsons, M. G. and R. L. Scott (2004). Formulation of multicriterion design optimization problems for solution with scalar numerical optimization methods. *Journal of Ship Research* 48(01), 61–76.
- Pepler, P. T. (2014). *The identification and application of common principal components*. Ph. D. thesis, Stellenbosch: Stellenbosch University.
- Saltelli, A., M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola (2008). *Global sensitivity analysis: the primer*. John Wiley & Sons.
- Sekitani, K. and Y. Yamamoto (1993). A recursive algorithm for finding the minimum norm point in a polytope and a pair of closest points in two polytopes. *Mathematical Programming* 61, 233–249.
- Spagnol, A., R. L. Riche, and S. D. Veiga (2019). Global sensitivity analysis for optimization with variable selection. *SIAM/ASA Journal on uncertainty quantification* 7(2), 417–443.
- Tanabe, R. and H. Ishibuchi (2020). An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing* 89, 106078.
- Trendafilov, N. and M. Gallo (2021). *Multivariate data analysis on matrix manifolds*. Springer.
- Trendafilov, N. T. (2010). Stepwise estimation of common principal components. *Computational Statistics & Data Analysis* 54(12), 3446–3457.
- Tukey, J. W. (1975). Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians, Vancouver, 1975*, Volume 2, pp. 523–531.

Wycoff, N., M. Binois, and S. M. Wild (2021). Sequential learning of active subspaces. *Journal of Computational and Graphical Statistics* 30(4), 1224–1237.

Zahm, O., P. G. Constantine, C. Prieur, and Y. M. Marzouk (2020). Gradient-based dimension reduction of multivariate vector-valued functions. *SIAM Journal on Scientific Computing* 42(1), A534–A558.

Zhang, Z., C. He, J. Ye, J. Xu, and L. Pan (2019). 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.

A Summary plots for the synthetic problem




Figure 4: Summary plots of SSPD, SEE, and AG for the synthetic problem applied to uniformly distributed data. AS1 stands for the active subspace of dimension 1, and f_1, f_2 are the two objectives.

B Impact of normalization on the RMSE performance

Table 3: Impact of the mean-variance normalization for a subset of problems. The values are the averages of the sum of RMSEs over 10 runs with the sample size $n = 1000$ in each of them. j stands for the dimensionality of the leading eigenspace: here are given the performances with respect to the two penultimate eigenspaces. The better performances are shown in bold.

(a) Normalized car side-impact

Method	j=5	j=6
AG	3.688	2.920
LP	4.712	3.872
SSPD	1.266	0.660
SEE	1.680	0.770
FG	1.312	0.764

(b) Original car side-impact

Method	j=5	j=6
AG	3.128	2.262
LP	3.260	2.329
SSPD	1.974	1.549
SEE	1.682	0.834
FG	2.012	1.574

(c) Normalized penicillin production

Method	j=5	j=6
AG	0.556	0.327
LP	0.748	0.583
SSPD	0.467	0.336
SEE	0.615	0.322
FG	0.452	0.353

(d) Original penicillin production

Method	j=5	j=6
AG	0.779	0.321
LP	0.777	0.220
SSPD	0.764	0.211
SEE	0.655	0.280
FG	0.717	0.193

(e) Normalized switching ripple

Method	j=6	j=7
AG	0.036	0.025
LP	0.038	0.024
SSPD	0.028	0.018
SEE	0.031	0.023
FG	0.028	0.019

(f) Original switching ripple

Method	j=6	j=7
AG	0.032	0.029
LP	0.034	0.029
SSPD	0.031	0.028
SEE	0.031	0.021
FG	0.031	0.028

C Local active subspace based on data depth

Given a data \mathbf{D}_n , a data depth $D(\mathbf{x}|\mathbf{D}_n) : \mathbb{R}^d \rightarrow [0, 1]$ measures how central (or deep) a point is in this data cloud Tukey (1975). It attains its maximum value for the deepest point in the cloud (for instance, the center of a symmetric distribution) and tends to zero for outlying points, i.e., it orders data by their degree of centrality. It possesses useful properties such as translation invariance $D(\mathbf{x} + \mathbf{b}|\mathbf{D}_n + \mathbf{b}) = D(\mathbf{x}|\mathbf{D}_n)$ for any $\mathbf{b} \in \mathbb{R}^d$, linear invariance $D(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{D}_n) = D(\mathbf{x}|\mathbf{D}_n)$ for every nonsingular $d \times d$ matrix \mathbf{A} and so on Mosler (2013).

From the perspective of this work, we observe that depth functions vary the most on the areas where the sample points are the denser, and vary much less in the sparser parts: this is in opposition with the objective-wise gradients, since large gradients of the depths may correspond to the areas where the functions vary less (and hence cluster), unless the function is multimodal. One can alleviate this problem by applying a depth function in the copula or rank space which amounts to applying the inverse cumulative distribution function of every objective, making the depth value invariant to any increasing monotonous transformation of the objectives. As a result it puts more weight on the parts where the functions change the most as illustrated in Figure 5, and thus it may be used as a weight function to favor the given parts of the image space of the function of interest.




Figure 5: Comparison of the zonoid Mosler (2013) and potential depths function on a test problem, in the original output space and in copula (rank) space.