Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

I'll tell you what I want: Categorization of Pareto Fronts for Automated Rule-based Decision-Making

Matthias Hoffmann, Thomas Schmitt, Kathrin Flaßkamp

2022

Preprint:

Copyright Elsevier 2022. This manuscript version is made available under the CC-BY-NC-ND 4.0 license http://creativecommons.org/licenses/by-nc-nd/4.0/

Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 55-16 (2022) 376–381

I'll tell you what I want: Categorization of Pareto Fronts for Automated Rule-based Decision-Making ⋆

Matthias K. Hoffmann * **Thomas Schmitt** **
Kathrin Flaßkamp *

* *Systems Modeling and Simulation, Saarland University, Saarbrücken, Germany (e-mail: {matthias.hoffmann,kathrin.flasskamp}@uni-saarland.de)*
** *Control Methods & Robotics Lab, Technical University of Darmstadt, Darmstadt, Germany (e-mail: thomas.schmitt@rmr.tu-darmstadt.de)*

Abstract: The application of Pareto optimization in control engineering requires decision-making as a downstream step since one solution has to be selected from the set of computed Pareto optimal points. Economic Model Predictive Control (MPC) requires repeated optimization and, in multi-objective optimization problems, selection of Pareto optimal points at every time step. Thus, designing an automated selection strategy is favorable. However, it is challenging to come up with a measure – possibly based on a Pareto front analysis – that characterizes preferred Pareto optimal points uniformly across different Pareto fronts. In this work, we first discuss these difficulties for application within MPC and then suggest a solution based on unsupervised machine learning methods. The approach is based on categorizing Pareto fronts as an intermediate step. This allows generating an individual set of rules for every category. Thereby, the human decision-maker's preferences can be modeled more accurately and the selection of a Pareto optimal solution becomes less time-consuming while breaking down the decision-making process into a selection solely based on the Pareto front's shape. Here, the measures act as anchor points for the decision rules. Lastly, a novel knee point measure, i.e. an approximation of the Pareto front's curvature, is presented and used for a knee point-focused categorization. The proposed algorithm is successfully applied to a case study for an energy management system. Moreover, we compare our method to using singular measures for decision-making in order to show its higher flexibility leading to better performance of the controller.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/4.0/)

Keywords: Multi-objective Optimization, Model predictive Control, Decision-making, Energy Management Systems, Optimal Control

1. INTRODUCTION

Multi-objective Optimization (MOO) provides a set of tools for dealing with multiple objectives whenever reformulation as a single-objective optimization problem is impossible or undesirable. Instead of having one optimal solution, MOO problems possess a continuum of optimal trade-offs between the objectives. The so-called human decision-maker (DM) has to select from this set the solution to be used within the considered application.

Model Predictive Control (MPC) is a control method that is widely used in applications, e.g. in the medical field (Czakó et al. (2017)), autonomous driving (Williams et al. (2018)), space travel (Pascucci et al. (2015)), energy management systems (Schmitt et al. (2020)), and many more. As a control method using optimal control, MPC can benefit from using MOO techniques, as there is typically more than one cost function interesting to an operator. When MOO is applied in this context, additional difficulties arise, since in MPC, similar optimal control problems, parameterized by time, have to be solved and solutions have to be applied repeatedly under real-time constraints. This makes MPC a dynamic MOO problem. The selection of a solution can happen manually, meaning the DM has to do this in every time-step, or by using measures that evaluate the Pareto front's shape automatically to select one solution. The first approach might be theoretically applicable in applications like energy management systems, where the time between two decisions is multiple minutes, but it becomes impossible for systems with shorter windows for decision-making. For this reason, measures for selecting points are applied. These make some assumptions on the Pareto front's shape, e.g. that it forms a knee. As the problem changes with time, so do the Pareto fronts and a measure-based selection produces unpredictable or even unwanted results. This work employs clustering techniques on Pareto fronts computed beforehand to help the DM with a better inclusion of his preferences. We promote a knee-based clustering and for that motivate a new measure that finds the point, where the Pareto front's radius of curvature is minimum.

⋆ Thomas Schmitt acknowledges the financial support by the Honda Research Institute Europe.

2405-8963 Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2022.09.053

1.1 Related Work

There is a large number of different methods for automatic selection of a Pareto optimal solution from a subset of the Pareto front. For once, there are methodologies to find compromise solutions, i.e. Pareto optimal points that are most similar or close to the ideal point or most dissimilar to the Nadir point. Examples of measures are the Euclidean distance to the ideal point or LINMAP by Srinivasan and Shocker (1973), that calculates a weighted distance. TOPSIS by Hwang and Yoon (1981) combines the shortest distance to the ideal point with the largest distance to the Nadir point.

Alternatively, knee points are considered to be special solutions to the MOO problem. Das (1999) points out, that human DMs typically select solutions in the middle of Pareto fronts where there is a bulge in the curve. Branke et al. (2004) propose two methods of finding these points from a Pareto subset. The angle-based focus calculates the reflex angle to each of the direct neighbors and selects the point where the reflex angle becomes maximum. In accordance with the idea of multiple weightings selecting the same Pareto optimal point, the authors also suggest a utility-based focus, where the points from the Pareto subset are ranked based on how many weights they could fulfill best if these were the only attainable solutions.

Schmitt et al. (2020) present the following measures that are applied to the normalized Pareto subset:

- closest to Utopia point (*CUP*) - the Euclidean distance to the Utopia point.
- angle to extreme points (*AEP*) - the angle spanned by connecting vectors from the Pareto optimal point to the extreme points.
- angle to neighbors (*ATN*) - the obtuse angle to the direct neighboring points, similar to the reflex angle.

1.2 Contribution and structure of this work

In this work, we automate the recurring decision-making process within an MPC loop by categorizing the occurring Pareto fronts. We assume that the prominence of a knee point is a major visual cue for human DMs, thus the categorization is based on factors that quantify the Pareto front's exhibition of knees via appropriately designed measures. As one of these measures, we propose the Pareto front's curvature, approximated at the computed Pareto optimal solutions. For each of the categories, the DM formulates rules to apply to Pareto fronts of that category, thus incorporating the preferences into the automatic selection – I'll tell you what I want. This algorithm breaks down the decision process by extracting the most common shapes, basing the decisions solely on these, reducing the load on the DM, and helping them build an algorithm imitating their decisions online.

This paper is organized as follows. In Section 2 we summarize the foundations of MOO. The idea of categorizing Pareto fronts is discussed in Section 3. For a knee-based categorization, this section also introduces an approximation of the radius of curvature as a new measure for selecting one Pareto optimal solution. For the example of an energy management system for a corporate building, the categorization is applied and its performance is evaluated in Section 4. It is shown that the selection of a Pareto optimal point based on the categorization yields preferable results, while providing more flexibility than singular metrics. Finally, a summary of this work is given in Section 5.

Notation: In this work, subindices are used to denote parts of known values, e.g. entries of a vector. Superindices are used to name specific values. For a list of parameter vectors, $\boldsymbol{p}^j$ marks the j-th parameter vector and p_i^j its i-th entry.

2. FOUNDATIONS

In this section we will give a brief overview on MOO, providing the techniques for dealing with k-dimensional cost functions.

The dynamic MOO problem (Azzouz et al. (2017)) is given by

$$\begin{aligned} \underset{\boldsymbol{p}}{\text{Minimize}}\ \boldsymbol{J}(\boldsymbol{p},t) &= [J_1(\boldsymbol{p},t), J_2(\boldsymbol{p},t), \dots, J_k(\boldsymbol{p},t)]^\intercal \\ \text{s.t. } g_j(\boldsymbol{p},t) &\le 0 \qquad j = 1,2,\dots,m \\ h_l(\boldsymbol{p},t) &= 0 \qquad l = 1,2,\dots,n, \end{aligned}$$

with the time t, the decision variables $\boldsymbol{p} \in \mathbb{R}^\nu$, k cost functions J_i, m inequality constraints g_j and the n equality constraints h_l. In the following, we will omit t, since the algorithms in this work do not depend on it. The constraints on $\boldsymbol{p}$ lead to a distinction between feasible and infeasible decision vectors. The set of all feasible decision vectors is then the feasible decision space $\mathbb{P} \subseteq \mathbb{R}^\nu$. Since the cost functions are dependent on the decision variables, the constraints on $\boldsymbol{p}$ result in constraints on $\boldsymbol{J}$ and with that the feasible criterion space is denoted by

$$\mathbb{J} = \{\boldsymbol{J}(\boldsymbol{p}) \mid \boldsymbol{p} \in \mathbb{P}\}.$$

When dealing with multiple cost functions, there typically does not exist a singular best solution. Instead, two solutions are compared with the concept of dominance. A solution $J(\boldsymbol{p}^*)$ is said to dominate another solution $J(\boldsymbol{p}')$, if $J_i(\boldsymbol{p}^*) \le J_i(\boldsymbol{p}')\ \forall i \in [1,k]$ and $\exists\, j \in [1,k]$ with $J_j(\boldsymbol{p}^*) < J_j(\boldsymbol{p}')$. The Pareto front is then the set of all non-dominated points in $\mathbb{J}$. In most applications, the Pareto front cannot be computed exactly. These are for example genetic algorithms or scalarization techniques. The latter transform the MOO problem to a scalar one and find one Pareto optimal point at a time. In this work, the term *Pareto subset* is used for $f = \{\boldsymbol{J}(\boldsymbol{p}_1), \boldsymbol{J}(\boldsymbol{p}_2), \dots\}$, a discrete set of Pareto optimal points.

The *ideal* point and the *Nadir* point describe the theoretically best and worst point in the criterion space in relation to the Pareto front. The minimal values of all objectives are collected to attain the unattainable ideal point

$$\boldsymbol{J}^{\text{ideal}} = [\min_{\boldsymbol{p}} J_1(\boldsymbol{p}), \min_{\boldsymbol{p}} J_2(\boldsymbol{p}), \dots, \min_{\boldsymbol{p}} J_k(\boldsymbol{p})],$$

whereas the Nadir point consists of the maximum values of all Pareto optimal points. Its entries are given by

$$J_i^{\text{nadir}} = \max_{\boldsymbol{p}} \{J_i(\boldsymbol{p}) \mid \boldsymbol{p} : \boldsymbol{J}(\boldsymbol{p}) \text{ is non-dominated } \}.$$

These points allow for a normalization of f to the range of $[0,1]$ in all objectives with

$$\tilde{J}_i(\boldsymbol{p}) = \frac{J_i(\boldsymbol{p}) - J_i^{\text{nadir}}}{J_i^{\text{ideal}} - J_i^{\text{nadir}}}.$$

The normalized Pareto subset is then given by $\tilde{f}$. Note that in this work normalized Pareto subsets are used.

Once a set of Pareto optimal points has been computed, one has to be selected according to the DM's preferences. Here, we distinguish between three types of decision-making.

(P1) A priori formulation - the DM knows how to formulate his preferences exactly, so a scalarization, e.g. weighted sum, can be applied.

(P2) A posteriori formulation - a set of Pareto optimal points is presented to the DM to select one of them.

(P3) No preference formulation - an algorithm selects a Pareto optimal point based on certain characteristics of the Pareto front.

For the second and third approaches, the Pareto front needs to be sampled. The a posteriori formulation requires the DM to select a solution in every iteration of the optimization. In the case of MPC, where the solution needs to be selected in execution time, this is not a feasible task. When no preference can be formulated, a measure that uses characteristics of the Pareto front, like the ones presented in Section 1.1, is applied to select a solution. For a profound review of MOO and more information on knee points we refer to the review by Marler and Arora (2004) and the work of Deb and Gupta (2011), respectively.

Various measures for automated decision-making (P3) can be defined, e.g. the reflex angle (cf. Section 1), which shall be maximized. In this work, we combine the second and third approach by finding and applying general rules that represent the DM's preferences better than singular measures could. In particular, we are interested in measures based on the existence of a knee point, i.e. reflex angle, *ATN*, or our *RoC*. The knee of the Pareto front is considered special, as for an improvement from these points in any objective, one has to accept an even larger deterioration in at least one other objective (Branke et al. (2004)). Sharp knees even correspond to multiple weightings, thus they have a higher probability of fulfilling the DM's preferences than other points, as shown in Fig. 1. However, if the shape

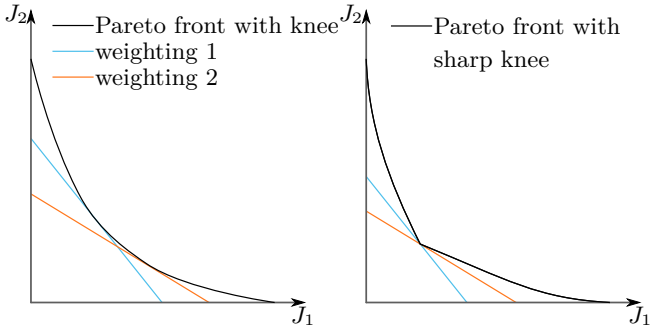


Fig. 1. The comparison of Pareto fronts with a "smooth" and a sharp knee is shown. A weighting of two cost functions can be interpreted as a line in the criterion space, their solution is the point closest to the ideal point, such that the line still touches the feasible criterion space. On a front with a smooth knee, different weights produce different, but close, solutions, whereas on fronts with a sharp knee multiple weights can correspond to the same solution.

changes as much as displayed in Fig. 2 and the selected point shifts to one extreme, this approach is not applicable anymore. Especially, when MOO is applied in MPC with vastly different initial states and external conditions, this can become a problem. More intelligent algorithms are needed to overcome this disadvantage, for example by selecting a useful measure fitting a Pareto front's shape.

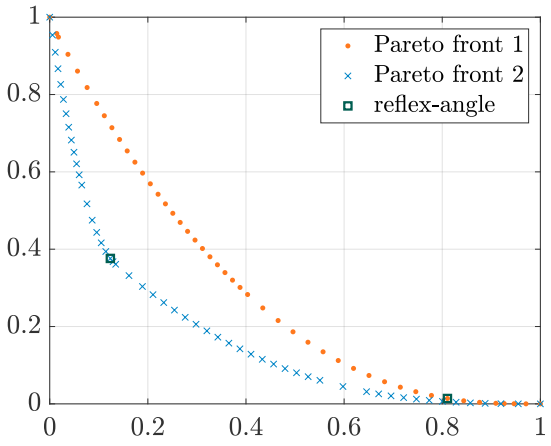


Fig. 2. Two dissimilar Pareto subsets of the same system dynamics with different time-dependent parameters are displayed, where the selection based on a singular measure cannot guarantee similar results.

3. CATEGORIZATION OF PARETO FRONTS

Without information other than the Pareto front's shape, it is reasonable to assume that the DM selects similar solutions for Pareto fronts with a similar shape. However, as Pareto fronts may differ in shape, particularly in knee prominence (cf. Section 2), the rules on which the DM selects solutions have to take into account these differences. In this section, we present a method for categorizing Pareto fronts based on the manifestation of a knee. We will first give a summary of the algorithm. After that, each step is explained in detail.

3.1 Summary

The categorization of Pareto fronts is done using clustering algorithms. First, a representative set $\mathbb{F} = \{\tilde{f}_i \mid i = 1,2,\ldots,N_f\}$ of N_f exemplary Pareto subsets $\tilde{f}_i$ is generated. Approximately 60 % of the Pareto subsets are collected to build up the set used for the actual categorization $\mathbb{F}_{\text{cat}} \subset \mathbb{F}$. The rest is accumulated in the set used for evaluation and rule generation $\mathbb{F}_{\text{eval}} = \mathbb{F} \setminus \mathbb{F}_{\text{cat}}$. This ensures independence of the data used in categorization and rule generation while enabling cross-validation.

Next, the data need to be prepared for categorization by computing a feature representation. This feature representation captures some properties of the Pareto front that are interesting to the DM and that the categorization will use to differentiate the Pareto subsets.

After generating the corresponding feature vector for each Pareto subset in $\mathbb{F}$, a clustering algorithm, e.g. k-means++, is applied to $\mathbb{F}_{\text{cat}}$ to find categories of Pareto subsets. Finally, the rules are generated by matching the Pareto subsets in $\mathbb{F}_{\text{eval}}$ with the categories and visualizing sets of matching categories.

The following steps concisely describe the algorithm.

(S1) Compute a representative set of normalized Pareto subsets $\mathbb{F} = \{\tilde{f}_1, \tilde{f}_2, \dots\}$.

(S2) Set a feature vector representation.

(S3) Divide $\mathbb{F}$ into $\mathbb{F}_{\text{cat}}$ and $\mathbb{F}_{\text{eval}}$ in a 60-to-40 ratio.

(S4) Compute the feature representation for all Pareto subsets in $\mathbb{F}$.

(S5) Apply a clustering algorithm, e.g. k-means++, to the feature representation in $\mathbb{F}_{\text{cat}}$ to get μ clusters C_i.

(S6) Match the fronts in $\mathbb{F}_{\text{eval}}$ with the clusters C_i.

(S7) For each C_i, visualize Pareto subsets from $\mathbb{F}_{\text{eval}}$ and set rules for each cluster.

Now, we will dive deeper into each step of the algorithm.

3.2 Representative set of normalized Pareto subsets

As this algorithm is designed to take into account strongly differing Pareto subsets, generating a representative set of normalized Pareto subsets is an important preliminary step. In the case of MPC, one can run multiple long simulations with a random selection of Pareto optimal points or generate single scenarios randomly.

3.3 Feature vector representation

The feature vector representation is comprised of the visual cues interesting to a DM on Pareto fronts.

Radius of Curvature - RoC As already mentioned, knee points are interesting in the decision-making process, as they can capture multiple preferences at once. Selection via one singular measure, like the reflex angle or *ATN*, can identify the point that is most suitable as a knee point, but not whether this point is actually preferable, as shown in Fig. 2. If a point from the Pareto subsets optimizes multiple measures with different properties, one can at least assume that this point is part of an actual knee. For that reason, multiple features evaluating a point's aptitude have to be applied for a knee point-focused categorization.

Here, we motivate a new measure called the *RoC* that approximates the radius of curvature of a Pareto front based on the discrete Pareto optimal points in the corresponding Pareto subset. First, we define the *RoC* in two dimensions. Then, we show that the given definition can be extended to arbitrary dimensions.

Fig. 3 shows the general idea of how the *RoC* finds knee points. A circle in two dimensions can be defined by three points. If these three points are relatively close to each other, this can be used to approximate the curvature of a given curve, as shown in Fig. 3a. If this relative closeness is not fulfilled, like in the case of a sharp knee point, the approximation becomes worse, see Fig. 3b, where the radius should be 0. In the case of Pareto fronts, this does not pose a problem, as the knee point still minimizes the measure of the *RoC*.

Let $\tilde{\boldsymbol{J}}(\boldsymbol{p}) \in \mathbb{R}^2$ be a normalized Pareto optimal point for which we want to calculate the aptitude of being a knee point. Further, let $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}}) \in \tilde{f}$ and $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}}) \in \tilde{f}$ be the points with the next lower and the next higher J_1-value to $\tilde{\boldsymbol{J}}(\boldsymbol{p}) \in \tilde{f}$, respectively. Finally, let $\boldsymbol{c}$ be the center of a circle connecting $\tilde{\boldsymbol{J}}(\boldsymbol{p})$, $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}})$ and $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}})$. Then, the distance from $\boldsymbol{c}$ to the other three points should be equal. This is given by

$$\left\|\boldsymbol{c} - \tilde{\boldsymbol{J}}(\boldsymbol{p})\right\|_2 = \left\|\boldsymbol{c} - \tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}})\right\|_2, \tag{1}$$

$$\left\|\boldsymbol{c} - \tilde{\boldsymbol{J}}(\boldsymbol{p})\right\|_2 = \left\|\boldsymbol{c} - \tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}})\right\|_2, \tag{2}$$

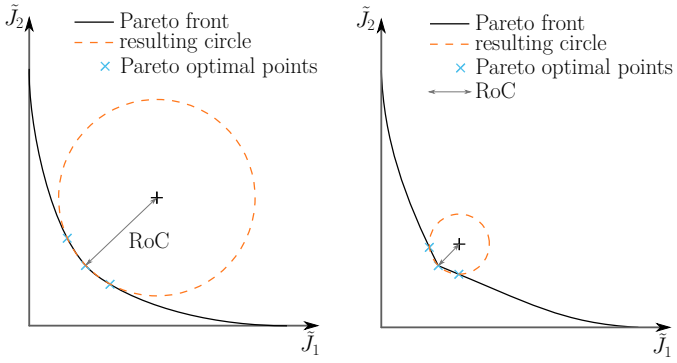
where $\|\cdot\|_2$ is the Euclidean norm. Equations (1) and (2) can be transformed to

$$2\begin{bmatrix}(\tilde{\boldsymbol{J}}(\boldsymbol{p}) - \tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}}))^{\intercal} \\ (\tilde{\boldsymbol{J}}(\boldsymbol{p}) - \tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}}))^{\intercal}\end{bmatrix}\boldsymbol{c} = \begin{bmatrix}\left\|\tilde{\boldsymbol{J}}(\boldsymbol{p})\right\|_2^2 - \left\|\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}})\right\|_2^2 \\ \left\|\tilde{\boldsymbol{J}}(\boldsymbol{p})\right\|_2^2 - \left\|\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}})\right\|_2^2\end{bmatrix}. \tag{3}$$

Under the assumption that $\tilde{\boldsymbol{J}}(\boldsymbol{p})$, $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{L}})$ and $\tilde{\boldsymbol{J}}(\boldsymbol{p}^{\text{H}})$ are not collinear, this system of equations can be solved for $\boldsymbol{c}$. With $\boldsymbol{c}$ calculated, the measure value of the *RoC* is given by

$$M_{\text{RoC}} = \left\|\boldsymbol{c} - \tilde{\boldsymbol{J}}(\boldsymbol{p})\right\|_2.$$

If the points are collinear, the *RoC* is defined to be infinity. The same applies to the extreme points. Comparison of $\boldsymbol{c}$ and $\tilde{\boldsymbol{J}}(\boldsymbol{p})$ allows for determining whether the bulge is convex ($\tilde{\boldsymbol{J}}(\boldsymbol{p})$ dominates $\boldsymbol{c}$) or concave ($\boldsymbol{c}$ dominates $\tilde{\boldsymbol{J}}(\boldsymbol{p})$).



(a) *RoC* on a discrete front without a sharp knee point, meaning a point that can be found with multiple weights.

(b) *RoC* on a discrete front with sharp knee point.

Fig. 3. The application of the *RoC* is shown on the fronts from Fig. 1. It approximates the radius of curvature in the middle point. This approximation is better for fronts without a sharp knee point, but is still minimal for sharp knee points on an evenly sampled Pareto front.

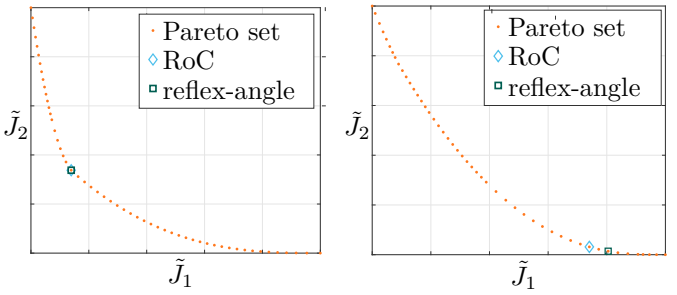


Fig. 4. Left: On Pareto fronts with a prominent knee, both measures (*RoC* and reflex angle) select the same point. Right: If the knee is only weak, the measures select different points. In this case a different measure like *CUP* could be more appropriate.

If more than two objectives are handled, (3) can be extended by more neighboring points. From our experience,

choosing the closest neighbors in each coordinate direction separately gives the best results. It is also possible to use all direct neighbors and find $\boldsymbol{c}$ using the Moore–Penrose inverse or one employs a weighted least squares approach in which the reciprocal Euclidean distance to each neighbor is used as the weights.

The *RoC* is very suitable for the following categorization of Pareto fronts. Its performance on Pareto fronts with a prominent knee is similar to the reflex angle and *ATN* while differing on Pareto fronts without a prominent knee, as shown in Fig. 4. For a knee-based categorization, we suggest using the following four features: minimum values of *ATN* and *RoC* for a Pareto subset and the position of the points minimizing these, rendering the feature vector $(2+2k)$-dimensional.

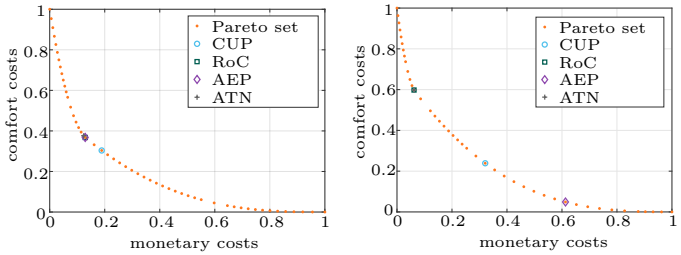
3.4 Clustering and definition of rules

After the feature representations of all Pareto subsets from $\mathbb{F}_{\text{cat}}$ are computed, they need to be clustered in order to find categories of Pareto subsets. There is a large number of clustering algorithms. Arguably, the most common one is k-means (Brunton and Kutz (2019)), but there are also others like mean shift (Cheng (1995)). When the cluster centers of the Pareto subsets in $\mathbb{F}_{\text{cat}}$ are found, the sets in $\mathbb{F}_{\text{eval}}$ need to be matched with these cluster centers using distance measures, e.g. the Euclidean distance.

The rules for each category are found visually. The DM examines a multitude of Pareto subsets from the same category and formulates rules they want to apply to similar sets, using the measures as anchor points for the rules. Rules are for example always selecting the *RoC* on Pareto fronts with prominent knee or the point best fulfilling a certain weighting.

4. RESULTS

The following simulations were done in MATLAB using the PARODIS framework (Schmitt et al. (2021)). The methods described in this work were applied to the MPC controlling the energy management system of a corporate building (see Schmitt et al. (2020)). Electrical power is drawn from the German public energy grid (P_{grid}), a photovoltaic plant (P_{ren}), and a combined heat and power plant (CHP, P_{chp}) and is fed to the energy consumers (P_{dem}). The



(a) If *RoC*, *ATN* and *AEP* coincide, this point is selected.

(b) Here, the *AEP* does not find the same point as *RoC* and *ATN*. The point closest to the midpoint of *RoC* and *CUP* is selected.

Fig. 5. Pareto fronts where the knee lies close to the ordinate are collected in this category. The *AEP* measure can be used to define the rule on which a DM selects since it deviates from the other measures if the knee point is extremely close to one extreme point.

thermal energy produced by the CHP ($\dot{Q}_{\text{chp}}$) is fed to the heating circuit and adds to the thermal energy produced by a natural gas heating ($\dot{Q}_{\text{heat}}$). An air-conditioner allows for cooling down the building ($\dot{Q}_{\text{cool}}$). Heat losses to the ground are denoted with $\dot{Q}_{\text{other}}$. A battery for storing energy is modeled with E. Using its stored energy and the building's temperature (ϑ_{b}) as states, the state-space representation is given by

$$\begin{bmatrix} \dot{E}(t) \\ \dot{\vartheta}_{\text{b}}(t) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & -\frac{H_{\text{air}}}{C_{\text{th}}} \end{bmatrix} \cdot \begin{bmatrix} E(t) \\ \vartheta_{\text{b}}(t) \end{bmatrix} + \begin{bmatrix} 1 & 1 & 0 & \frac{1}{\varepsilon_{\text{c}}} \\ 0 & \frac{1}{c_{\text{CHP}} \cdot C_{\text{th}}} & \frac{1}{C_{\text{th}}} & \frac{1}{C_{\text{th}}} \end{bmatrix} \cdot \begin{bmatrix} P_{\text{grid}} \\ P_{\text{chp}} \\ \dot{Q}_{\text{heat}} \\ \dot{Q}_{\text{cool}} \end{bmatrix} + \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & \frac{H_{\text{air}}}{C_{\text{th}}} & \frac{1}{C_{\text{th}}} \end{bmatrix} \cdot \begin{bmatrix} P_{\text{ren}} \\ P_{\text{dem}} \\ \vartheta_{\text{air}} \\ \dot{Q}_{\text{other}} \end{bmatrix},$$

where ϑ_{air} denotes the outer temperature. The system is discretized using a step size $\Delta T = 0.5s$.

As for cost functions, the monetary costs and the so-called comfort costs were analyzed. The monetary costs include the German industrial energy pricing for consumption and production, including a penalty for the highest consumption peak, and the costs for operating the CHP and gas heating (Schmitt et al. (2020)). Comfort costs are the quadratic deviation of ϑ_{b} to a set temperature, in this case 21°C. For the categorization, 17521 Pareto fronts, the equivalent of a one-year simulation, were computed.

4.1 Knee-based characterization of Pareto fronts

As mentioned in Section 3, the proposed feature vector for knee-based categorization is selected. The Pareto fronts are categorized using the k-means++ algorithm with the squared euclidean distance metric, as it is implemented in MATLAB. k is found empirically to be 6. After determining the categories and applying the categorization to the Pareto fronts of the evaluation set, Pareto fronts from the evaluation set belonging to the same category are visualized. Fig. 5 shows two Pareto fronts of the same category. This category collects fronts, where the Pareto front exerts a knee point that is located very close to the ordinate. When visualizing these fronts, two general cases can be distinguished and mathematically formulated. If *RoC*, *ATN*, and *AEP* coincide, the corresponding solution is selected. If the *AEP* selects a different point, this corresponds to an extreme solution for the knee point. In order to prevent high deviations of the set temperature, i.e. deviations in the comfort cost criterion, the Pareto optimal point closest to the midpoint of *CUP* and *RoC* is selected. The second part of this decision rule is just one possibility of how the DM can influence the decision process. By always selecting the knee point, the monetary expenses can be decreased, at the cost of increasing the temperature deviation.

4.2 Application of knee-based rules

Fig. 6 shows the total costs of one-year simulations of the energy management system, controlled with economic MPC. Six simulations were conducted, where the selection of a Pareto optimal point was performed by one of the presented measures, *CUP*, *AEP*, *ATN*, *RoC*, and two

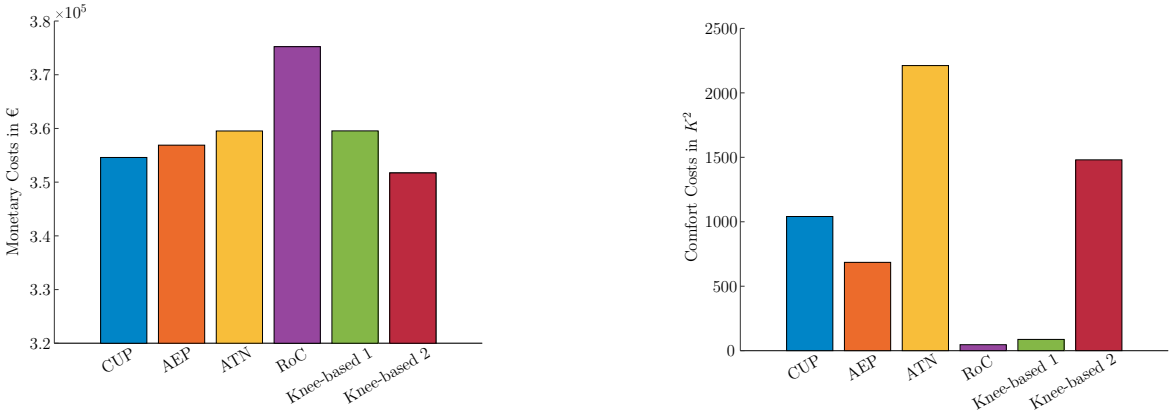


Fig. 6. Summed up costs for the one-year simulation of an energy management system controlled by an MPC. The selection uses either one measure or one of two knee-based clustering approaches. One of the clustering approaches achieves very low comfort costs, while also ensuring relatively low comfort costs, whereas the other is able to reduce the costs below the level of the *CUP*, by allowing higher comfort costs.

different knee-based approaches. The first knee-based approach (green) selects knees, if the knee point is not too close to one extreme point (cf. Fig. 5). This approach cuts, compared to the *RoC*, the monetary costs to the level of the *ATN*, while reducing the comfort costs close to the lowest found solution, the *RoC*. The second approach (dark red) focuses more on the monetary costs, reducing them below the ones produced by selection using *CUP*. Note that only *CUP*, *AEP*, *ATN*, and *RoC* were used to construct the rules. The categorization allows to even step out of their singular performances, proving the rules' high flexibility.

5. SUMMARY

In this work, we introduced the idea of categorizing Pareto fronts and setting rules based on these categories for better inclusion of the decision-maker's preferences. The algorithm clusters offline computed Pareto fronts, helping a human decision-maker to build up a set of rules to imitate his decision process online. This allows for adaptive application of measures or even their combination based on the Pareto front's shape. For this, the *RoC*, a new knee point measure approximating the Pareto front's radius of curvature, was introduced to perform a knee point-based categorization in combination with the angle to the nearest neighbors. The success of this approach was shown in a case study of an economic model predictive controller controlling an energy management system. Here, we compared two sets of rules based on the categorization approach with the singular measures showing our approach's flexibility as well as its improved performance in automated decision making.

REFERENCES

Azzouz, R., Bechikh, S., and Said, L.B. (2017). Dynamic multi-objective optimization using evolutionary algorithms: a survey. In *Recent advances in evolutionary multi-objective optimization*, 31–70. Springer.

Branke, J., Deb, K., Dierolf, H., and Osswald, M. (2004). Finding knees in multi-objective optimization. In *International conference on parallel problem solving from nature*, 722–731. Springer.

Brunton, S.L. and Kutz, J.N. (2019). *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press.

Cheng, Y. (1995). Mean shift, mode seeking, and clustering. *IEEE transactions on pattern analysis and machine intelligence*, 17(8), 790–799.

Czakó, B., Sápi, J., and Kovács, L. (2017). Model-based optimal control method for cancer treatment using model predictive control and robust fixed point method. In *2017 IEEE 21st International Conference on Intelligent Engineering Systems (INES)*, 000271–000276. IEEE.

Das, I. (1999). On characterizing the "knee" of the pareto curve based on normal-boundary intersection. *Structural optimization*, 18(2-3), 107–115.

Deb, K. and Gupta, S. (2011). Understanding knee points in bicriteria problems and their implications as preferred solution principles. *Engineering optimization*, 43(11), 1175–1204.

Hwang, C.L. and Yoon, K. (1981). Methods for multiple attribute decision making. In *Multiple attribute decision making*, 58–191. Springer.

Marler, R.T. and Arora, J.S. (2004). Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6), 369–395. doi:10.1007/s00158-003-0368-6.

Pascucci, C.A., Bennani, S., and Bemporad, A. (2015). Model predictive control for powered descent guidance and control. In *2015 European Control Conference (ECC)*, 1388–1393. IEEE.

Schmitt, T., Engel, J., Hoffmann, M., and Rodemann, T. (2021). Parodis: One MPC framework to control them all. Almost. *2021 IEEE Conference on Control Technology and Applications (CCTA)*, 466–471.

Schmitt, T., Rodemann, T., and Adamy, J. (2020). Multi-objective model predictive control for microgrids. *at-Automatisierungstechnik*, 68(8), 687–702.

Srinivasan, V. and Shocker, A.D. (1973). Linear programming techniques for multidimensional analysis of preferences. *Psychometrika*, 38(3), 337–369.

Williams, G., Drews, P., Goldfain, B., Rehg, J.M., and Theodorou, E.A. (2018). Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6), 1603–1622.