DOI: 10.1111/cgf.13990
Eurographics Conference on Visualization (EuroVis) 2020
M. Gleicher, T. Landesberger von Antburg, and I. Viola
(Guest Editors)

Volume 39 (2020), Number 3

PAVED: Pareto Front Visualization for Engineering Design

Lena Cibulski [1], Hubert Mitterhofer [2], Thorsten May [1] and Jörn Kohlhammer [1,3]

[1] Fraunhofer IGD, Darmstadt, Germany
[2] Linz Center of Mechatronics GmbH, Linz, Austria
[3] Technical University of Darmstadt, Darmstadt, Germany

Abstract

Design problems in engineering typically involve a large solution space and several potentially conflicting criteria. Selecting a compromise solution is often supported by optimization algorithms that compute hundreds of Pareto-optimal solutions, thus informing a decision by the engineer. However, the complexity of evaluating and comparing alternatives increases with the number of criteria that need to be considered at the same time. We present a design study on Pareto front visualization to support engineers in applying their expertise and subjective preferences for selection of the most-preferred solution. We provide a characterization of data and tasks from the parametric design of electric motors. The requirements identified were the basis for our development of PAVED*, an interactive parallel coordinates visualization for exploration of multi-criteria alternatives. We reflect on our user-centered design process that included iterative refinement with real data in close collaboration with a domain expert as well as a summative evaluation in the field. The results suggest a high usability of our visualization as part of a real-world engineering design workflow. Our lessons learned can serve as guidance to future visualization developers targeting multi-criteria optimization problems in engineering design or alternative domains.*

CCS Concepts
• ***Human-centered computing*** → *Visual analytics;* • ***Applied computing*** → *Engineering;*

1. Introduction

The world we live in today would be inconceivable without the outcomes of engineering. It is at the basis of systems, machines, and processes that drive many diverse areas of our society. We refer to the design stage of the engineering process as *engineering design*. At their core, most real-world decisions pose a conflict between several criteria. Engineering design is no exception: the key challenge is a *multi-criteria optimization* problem. Designs need to satisfy conflicting requirements: efficiency costs money, safety adds to complexity, durability increases material demands. Added to these requirements are various performance indicators like stress, deformation, or heat dissipation. A unique optimal solution does generally not exist. Instead, engineers must find a trade-off in a typically large solution space. Rather than the perfect solution, they search for a solution that is optimal for a given application.

To be utilizable by an optimization algorithm, preferences need to be made explicit in an objective function (a priori). This is often not desired, because preferences are typically vague and evolve over time. Without an objective function, optimization algorithms can only compute a number of mathematically equally good solutions, known as the *Pareto front*. From there, it is the responsibility of the engineer to choose the solution to be realized as a prototype. This can be classified as *a-posteriori* decision-making [Hor96].

With simulation being extensively used in engineering design, many optimization problems typically involve thousands of alternative solutions. Large Pareto fronts, multi-criteria solution spaces, and unstable subjective preferences make the decision process challenging. Engineers need assistance in making sense of the available options. Visualizing the Pareto front can help make informed decisions by highlighting conflicts and trade-offs, conveying their nature, and making the effects of applied preferences visible.

In this paper, we describe a design study addressing the role of multivariate Pareto front visualization for decision-making in engineering design applications. The collaboration with domain experts over a period of 1.5 years focused on the design and optimization of electric motors. Inspired by Sedlmair et al.'s nine-stage framework [SMM12] and Munzner's nested model [Mun09], we reflect on the user-centered design of a visualization tool that supports the process of selecting the most-preferred solution from a large set of Pareto-optimal designs. Conducting the design study yielded an interesting set of insights concerning two main topics: 1) the engineers' expectations and acceptance regarding visualization, and 2) the collaborative aspects of the user-centered visualization design. Our lessons learned from working with engineers to examine the potential of visualization for their tasks offer guidance to those who seek to develop visualizations for design engineers in the future.

In particular, the contributions of this paper include:

- A characterization and abstraction of the data, tasks, and requirements related to engineering design
- The design and evaluation of an interactive visualization of alternatives in multi-criteria optimization problems
- Reflections on the design process and common guidelines

2. Domain Characterization and Abstraction

We start by developing a characterization of the targeted problem in the field of electromechanical engineering. First, we provide an introduction to the design and optimization of electric motors (Section 2.1). Based upon that, we provide an abstraction of the data and tasks that engineers face when searching for the most-preferred design (Section 2.2). From these abstractions, we derive design requirements to be addressed by the visual design (Section 2.3).

2.1. Electric Motor Design Background

Electric motors have become an indispensable part of many industrial and domestic applications, from microdrives in electric toothbrushes to high-performance motors in transportation systems. In 2013, about 70% of the electrical energy in industry was consumed by electric motors [ZBL*13]. Their performance thus affects key indicators like energy consumption or productivity of the driven process. Additional requirements can involve fault tolerance, good controllability, compactness, and cost-efficiency. This places high demands on the design and optimization of electric motors.

Design engineers specify the geometry, material, winding patterns, etc. of an electric motor such that its performance and overall properties optimize given requirements. Up to a dozen of these design parameters are usually considered in the optimization process [DI13]. The evaluation of a motor's operational behavior is realized using simulation. An optimization algorithm, typically population-based methods like genetic algorithms, then computes a set of Pareto-optimal solutions [ZBL*13]. After validation of the results, the engineer chooses the most-preferred compromise. This selection is usually verified by additional simulations or experimental validation before the corresponding motor is taken to production.

Commercial tools for the design of electric motors provide only two-dimensional Pareto front visualizations that are not suited for optimization with multiple criteria. Therefore, our collaborators use their own optimization tool called *SyMSpace* [SKW*18] (formerly *MagOpt* [SBD*16]). Visual inspection of the Pareto front is performed using an interactive scatterplot matrix conveying pairs of the criteria to be optimized. The motor experts are quite familiar with concepts like brushing and linking. A selection of alternatives can be created, refined, and observed in linked histograms showing the related design parameters. Still, the analysis in SyMSpace is limited to two-dimensional projections of the Pareto front.

One challenge for the choice of a solution is the large number of available options, as a Pareto front can easily contain 100 to 200 multidimensional alternatives. Another challenge is the handling of conflicts between criteria, in particular when applying constraints. Due to the manufacturing tolerances to be expected during production, the selected solution also needs to be tolerant towards slight design parameter changes. The engineer's primary needs can be summarized as: a simultaneous overview of both criteria and alternatives together with an efficient drill down, perception of redundancies and conflicts among criteria, and sensitivity analysis.

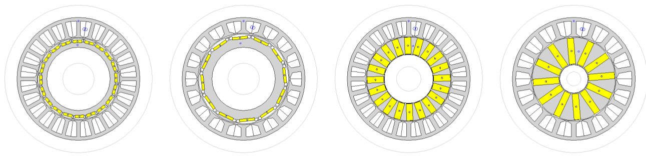

(a) *Neodymium, 20 magnets* **(b)** *Neodymium, 12 magnets* **(c)** *Ferrite, 20 magnets* **(d)** *Ferrite, 12 magnets*

Figure 1: *Four topologies of an internal rotor design resulting from combinations of magnet material and number of magnets. Variation of geometry, winding patterns etc. yields the actual design options.*

The design and optimization of motors is mostly conducted as commissioned work. This introduces a second type of stakeholders: the engineers' customers. Not all customers are experts in motor development themselves. Depending on their level of experience, either high trust is put in the choice made or the engineer is asked for clarification about design decisions. Our primary target users are the engineers responsible for the design of electric motors. The support of a joint decision-making between engineers and their customers will be addressed in future research.

2.2. Data and Task Analysis

This section covers the data, tasks, and design requirements for this design study, moving from domain-specific details to abstractions.

PMSM Simulation Model

The simulation considered as a running example throughout this work describes the operational behavior of a permanent magnet synchronous motor (PMSM) whose intended function is to drive a fan that cools the engine of a vehicle. The order made by an automotive supplier contains several specifications to be met: a rated power of 700 watt, a rated torque of 2.6 Newton meters, an outer diameter smaller than 136 millimeters, and an internal rotor. The motor should also fit the existing system setup in terms of size and shape. The customer's major interests are power and cost efficiency, small length, smooth running, and simple power electronics.

The design engineer has narrowed down the design space to either ferrite or neodymium for the *magnet material* as well as 12 or 20 for the *number of magnets*. All four combinations make up the available motor *topologies* (Figure 1). Within each topology, between eight and ten parameters related to geometrical dimensions, winding patterns, or material properties are varied stepwise. Any combination of parameters is called a *design option*. For each design option, the simulation evaluates the motor's operational behavior in terms of the criteria stator length, costs, power loss, maximum current, and torque ripple. The optimization returns 359 design options that are Pareto-optimal with respect to each topology considered separately. Options that are geometrically invalid or do not meet specified hard constraints are excluded during this process.

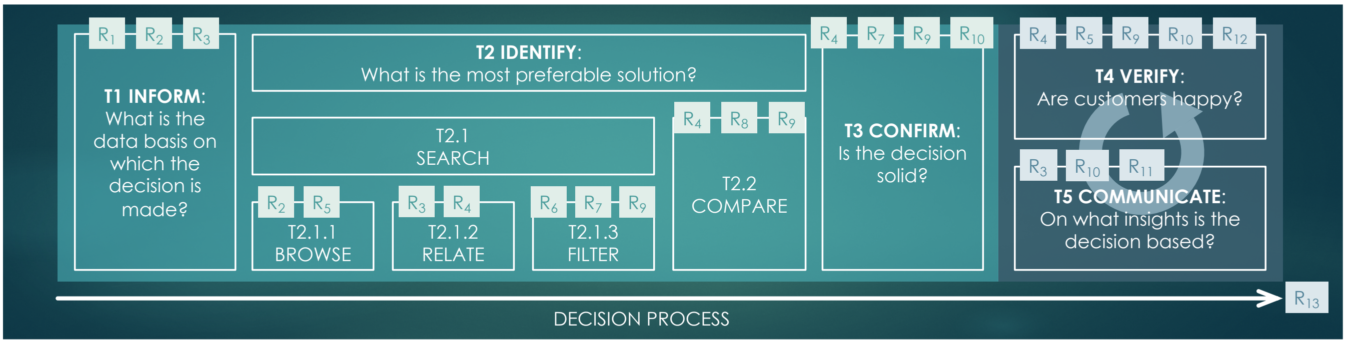

Figure 2: *An abstraction of the analysis tasks for multi-attribute choice in engineering design. T1 to T3 reflect the tasks that the engineer faces. T4 and T5 involve the customer as additional stakeholder. Derived design requirements R_1 to R_{13} are assigned to the respective tasks.*

Data Abstraction Simulation models are basically input-output models that approximate a function $X \to Y$ mapping some input dimensions $X = \{X_1, ..., X_n\}$ to a number of output dimensions $Y = \{Y_1, ..., Y_m\}$. We refer to the input dimensions X as *design parameters*. The dependent output dimensions Y are known as *criteria*. For the exemplary PMSM model, $n \in \{8, 9, 10\}$ and $m = 5$. Each criterion needs to be either minimized or maximized. The information about the desired direction of change is given as meta-data. The union $(\vec{x}, \vec{y})$ of a design option $\vec{x} = (x_1, ..., x_n); x_i \in X_i$ and its criteria $\vec{y} = (y_1, ..., y_m); y_i \in Y_i$ as provided by a simulation run is called *alternative*. In contrast to the motor experts, who adopted the terms "objectives" and "individuals" from genetic optimization theory, we use the terms "criteria" and "alternatives", as they provide a better link to the field of multi-criteria decision-making.

The challenge of engineering design lies in the absence of a direct inverse relation $Y \to X$ [STDS95]. Different Pareto-optimal design options thus need to be explored, which are computed by an optimization algorithm based on regular sampling of the input space. The sampling range and step size is specified separately for each design parameter. The final Pareto front contains a few hundred alternatives where no criterion can be improved without sacrificing at least one other criterion. Our collaborators do not expect to need more than ten criteria to reflect their customers' interests.

Analysis Tasks and Workflow As Ullman states: *"[engineering] design is decision-making"* [Ull01]. The decision-maker's primary goal is to identify the solution that best matches their customers' interests within the specified hard constraints. Given the previous data abstraction, this goal refers to the task of *multi-attribute choice* as defined by Dimara et al. [DBD17]. It describes the identification of the best among a finite number of multi-criteria alternatives that are known ahead of time.

A variety of strategies for multi-attribute choice have been studied in decision theory. In their position paper, Torsney-Weir and colleagues propose to consider such strategies for visual tool design [TWSM15]. We therefore began the task characterization by classifying the decision strategy of our primary domain expert. According to the taxonomy described by Payne et al. [PPBJ93], the expert's decision-making process is most similar to the *elimination by aspects* strategy. This strategy is about filtering options into acceptable and unacceptable regions until a final choice remains.

Our task abstraction (Figure 2) is based on selected tasks from 12 taxonomies that we found in the visualization literature. Please refer to the supplementary material for details on the task extraction and analysis questions underlying the selected tasks. Our task abstraction is not specific to the engineering domain and can be mapped to any multi-criteria decision-making scenario.

The decision process starts with the *inform* task (T1). It includes an inspection of the optimization results for their validity to answer questions like "Does the simulation produce plausible results?". The task is also about gaining a first overview of the design space, i.e. "What is the shape of the Pareto front? How diverse are the alternatives?", as well as the criteria space, i.e. "What is the distribution of alternatives? What is the nature of conflicts?".

Next, the actual decision-making takes place: the decision-maker needs to identify the most-preferred alternative (T2). Ullman states that two thirds of activity spent on engineering design tasks is related to searching the design space [Ull01]. The *search* task (T2.1) includes sub-tasks like browsing through the alternatives (T2.1.1), developing preferences as relations between criteria become apparent (T2.1.2), and using these preferences to judge and filter alternatives (T2.1.3). The search phase typically results in a subset of interest and is followed by a comparison phase. The *compare* task (T2.2) is primarily about judging the superiority of alternatives with respect to the preferences developed throughout the search phase. Of course, the identification of the most-preferred alternative might involve going back and forth between the sub-tasks.

As the decision-making is characterized by conflicting criteria, analysts need to *confirm* (T3) their decisions to increase confidence in their choice. Confirmation includes reviewing the perceived quality of the chosen design, revisiting its superiority compared to other favorite designs, and checking its sensitivity with respect to minor changes of the design parameters.

Once the decision has been confirmed by the engineer, it needs to be presented to the customer. It is highly important that engineers *verify* (T4) the chosen design with their customers and *communicate* (T5) on what insights the decision is based. One domain expert summarized why this is so important: *"The decision-making process should be comprehensible for the customer to prove its plausibility"* (A3). This might also include general recommendations with respect to why certain options should be considered or avoided.

2.3. Design Requirements

From the aforementioned data and task abstractions, we have derived the following design requirements to guide the visual design:

R_1 *Validation* – Show criteria ranges for simulation steering
R_2 *Overview* – Provide an overview of complete alternatives such that any criteria value of any alternative can be retrieved
R_3 *Criteria Relations* – Highlight redundant or conflicting criteria
R_4 *Trade-offs* – Support subjective perception of superiority
R_5 *Filter* – Support perception of the effect of constraints
R_6 *User Interaction* – Support simple and effective selection
R_7 *Provenance* – Store favorite alternatives for future comparison
R_8 *Comparison* – Support criteria-wise comparison of alternatives
R_9 *Details* – Enable direct reading of raw criteria values
R_{10} *Sensitivity* – Show design parameter values of alternatives
R_{11} *Transparency* – Support awareness of intermediate decisions
R_{12} *Accessibility* – Make the decision accessible to customers
R_{13} *Export* – Provide the data of the selected design for production

3. Related Work

Previous design studies on multi-criteria decision-making have been conducted in domains other than engineering design. *SOMMOS* focused on decision-making as a high-level goal [CAS*13]. Although it is applicable to different domains, additional potential might be unlocked by considering domain-specific tasks and requirements. *Vismon* was introduced as a trade-off analysis tool for users working in fisheries management [BMPM12]. Arbesser et al. provide a high-level view on challenges and lessons learned from distributing visualizations to engineers [AMKP17]. *RelEx* resulted from a design study on optimization of traffic flow in communication networks [SFMB12]. They address an engineering design problem on a different data abstraction, namely graphs. Basole et al. present a visual analytics tool to support the design of complex engineered systems [BQP*15]. However, they focus on the early design phase and the modeling of a system, not on multi-attribute choice from a number of generated alternatives. To the best of our knowledge, no design study on multi-attribute choice exists that particularly addresses the needs of users from engineering design.

3.1. Multivariate Pareto Front Visualization

Depicting a Pareto front as a set of multi-criteria alternatives [KW08] requires visualization methods that effectively map large multivariate data to two-dimensional visual space [LM08]. Either dimension reduction or lossless projection [DBD17] can be used.

A popular approach to dimension reduction is the self-organizing map (SOM) [Koh01]. In *SOMMOS*, Chen et al. semantically enhance a SOM with criteria anchors on a regular convex polygon as well as radial bar charts depicting individual alternatives [CAS*13]. In engineering design, the SOM has been employed for criteria-wise design space exploration in the context of aerodynamic optimization [OJC05]. Zhao et al. use t-SNE as a projection method for their system *SkyLens* and also employ radial glyphs to represent individual alternatives [ZWC*17]. While useful for navigating Pareto fronts, dimension reduction does not allow for a direct retrieval of any criteria value from any alternative. Thus, we do not consider SOM or t-SNE as applicable projections for our purpose.

In contrast, a lossless projection preserves the raw information. Tabular visualizations have been extended with weight-based ranking to simplify multi-attribute choice [GLG*13, SOL*15]. To help users better understand the effects of different weights, *Weightlifter* has been introduced [PSTW*16]. However, preferences are often too complex to be captured by weights. Mühlbacher et al. use scatterplots to visualize the trade-off between accuracy and complexity of decision trees [MLMP17]. Scatterplot matrices depict more than two criteria, e.g. for the design of an aircraft engine [MGH*07]. Similar to our consideration of motor topologies, the authors observe distinct engine concepts. Lotov et al. depict ensembles of two-dimensional Pareto fronts that result from discretely varying a third criterion [LBK13]. Still, scatterplot-based methods provide a limited perception of complete multi-criteria alternatives.

For visualizing Pareto fronts with respect to all criteria simultaneously, parallel coordinates are predominantly used [BC03]. Andrienko and Andrienko propose modifications regarding axis orientation, scaling, alignment, and ranking of alternatives [AA01]. Among others, parallel coordinates have been applied to design problems from automotive engineering [MJJ*05, BPFG11], aerospace engineering [SKW98, GBS*99], and aerodynamic engineering [KIPS13]. Fleming et al. describe the use of parallel coordinates for multi-criteria optimization in real-world engineering design, but do not study the visualization design [FPL05]. Parallel coordinates used in engineering mainly depict the design space together with up to three criteria. Our approach provides support in interactively exploring trade-offs involving up to ten criteria, while also considering the design space.

3.2. Visual Parameter Space Analysis

Gaining insight into the correspondence between design options and performance is an important aspect of multi-attribute choice [SKW98]. The same goal is pursued by parameter space analysis. Sedlmair et al. provide a conceptual framework that includes a data flow model, navigation strategies, and analysis tasks [SHB*14]. The current body of work in visual exploration of parameter spaces mainly comprises approaches from different application areas like meteorology [PWB*09] or image analysis [PBCR11, TWSM*11]. Beham et al. employ parallel coordinates to relate the parameter space of a geometry generator to the resulting shapes [BHGK14].

Closest to our work are approaches addressing simulation-based engineering. Matkovic et al. use interactive visual analysis for exploring the parameter spaces of fuel injection systems [MJJ*05]. *ParaGlide* allows for interactive partitioning of parameter spaces exemplified with a fuel cell simulation [BSM*13]. Berger et al. propose an interactive approach for car engine design that enables users to navigate a multivariate design space while observing the behavior of multiple criteria [BPFG11]. Originating from a focal design, star sampling in the design space is performed to observe criteria ranges that are within reach. The sensitivity of criteria to changes of the focal point is conveyed by neighborhoods in the criteria space being mapped back into the design space.

Previous work related to specific visual encodings and interaction techniques is discussed in the context of the visual design in Section 5. Previous work on design study methodology and evaluation is referred to in the respective Sections 4, 6, and 7.

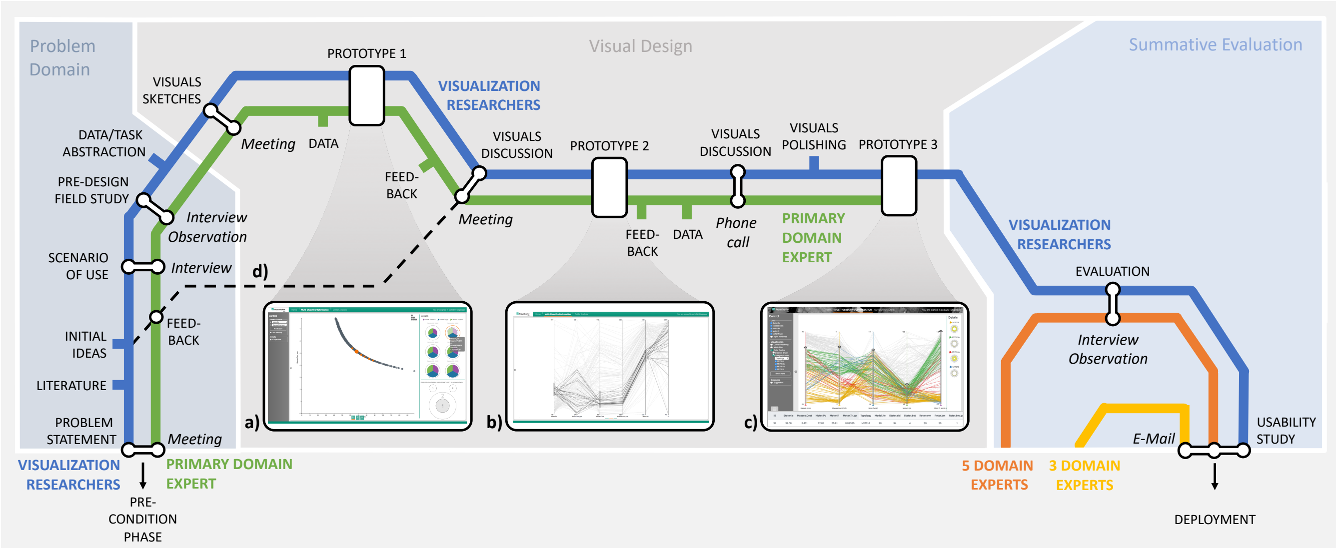

Figure 3: *An overview of our design process, inspired by storyline visualizations [TM12]. A mechatronics scientist (green) accompanied us (blue) during the domain characterization and visual design stage. The process involved two intermediate prototypes (a, b). The final prototype (c) was qualitatively evaluated by a group of five engineers (orange). For the usability testing, three additional engineers (yellow) were considered. A shortcut in the design phase (d) could have been taken if we had listened more carefully to the primary expert's feedback.*

4. Iterative Design Process

Our user-centered design process was organized in three stages: 1) characterization of the problem domain, 2) visual design, and 3) summative evaluation. An overview of the design process is presented in Figure 3. A mechatronics scientist at LCM (a co-author of this paper) with extensive experience in the design and optimization of electric motors was our primary domain expert. He accompanied the first two stages of the design process with constant insights into the domain on the one hand and feedback to our visual design on the other hand. The exchange took place in the form of 1) scheduled meetings in person, where fundamental characterization and design aspects were discussed, 2) phone calls for instant clarification of open issues, and 3) e-mail communication for summary feedback as well as confirmation of our documented insights. The primary outcomes of this process are the prototypes.

The understanding of the problem domain that led to its characterization was informed by different sources. We started by reading about the target domain background, in particular literature suggested by the domain expert. Asking the expert about tasks, tools, practices, and challenges in multiple sessions and discussing a scenario of use provided us with a fundamental understanding of the domain. A pre-design field study, where we observed the domain expert on a real-world use case, made sure that we did not run into the threat of mischaracterizing the problem [Mun09]. By having the primary expert constantly review the identified key tasks and their abstractions we ensured a common understanding of the use case and additionally contributed to an immediate validation.

The design stage mainly involved in-person meetings where we sketched and discussed visual encodings using the same pen and paper. To minimize the learning effort, we decided to start the development of our visual design from the interactive scatterplots that the domain experts already used for pairwise trade-offs. In multiple iterations, we implemented and refined an initial prototype, shown in Figure 3a. At its core, it contained an augmented scatterplot, where points were interactively shown as radial bar charts [CAS*13] to encode additional criteria. Solutions of interest could be cached and used for a detailed, criteria-wise comparison.

However, the expert's criticism of this visual encoding discouraged us from following up on the glyph-based scatterplot. It made us realize that we did not prioritize the criteria-wise overview high enough in our initial task abstraction. As a consequence, we discussed other visualization designs regarding the ability to convey both an alternative-wise and a criteria-wise overview. We then developed a second prototype that made use of parallel coordinates (Figure 3b) and iteratively developed it towards a parallel coordinates visualization, whose interactions better reflect the optimization operations performed by the engineer (Figure 3c). In retrospect, the detour via the first glyph-based prototype would not have been necessary, if we had listened more carefully to the domain expert, who already expressed his interest in parallel coordinates early in the domain characterization process (Figure 3d).

We performed a downstream validation against the threats of problem mischaracterization and wrong abstractions [Mun09]. For this, we evaluated the usefulness of the final tool with a qualitative field study and a quantitative usability scale. Both were conducted with five domain experts other than the primary domain expert. The usability study was extended by an additional group of three domain experts. We also validated the domain characterization by following up on the experts' interest in adopting the tool for their daily work. The evaluation details are discussed in Section 6.

5. PAVED Design

Motivated by the domain characterization and abstraction, we designed the visual encodings and interactions of PAVED, a parallel coordinates visualization that supports analysts in exploring Pareto-optimal designs to make an informed preferential choice. While it has been designed for engineers, our approach generalizes to any multi-criteria decision where the preferences cannot be quantified in advance. It is our ambition to provide the simplest solution that works well for the described multi-criteria optimization problem. Our visualization is designed to be intuitive, easy to learn and seamlessly integrated into the engineers' workflow. Simplicity is achieved by a one-view approach and a reduced yet effective user interaction for drilldown. To not neglect potentially relevant information, the focus is on simultaneously depicting all design options with all associated criteria (R_2 Overview). Raw data can be accessed at any time in a tabular view (R_9 Details, R_{13} Export).

5.1. Design Rationales

Before we present the actual visual encodings, we describe high-level design rationales that result from the design requirements.

Prefer simple over flexible user interaction As elimination-by-aspects is the prevalent strategy, the decision process is all about eliminating undesired alternatives (R_6 User Interaction) to move towards a small subset of favorites, from which the final choice is made. Thus, interaction should not demand more effort than absolutely necessary to achieve an intended selection. This means to reduce the interaction to the minimal set of operations needed for the fundamental optimization tasks. 'Simple' also includes that a selection is easy to describe, e.g. by means of a range, to effectively communicate decisive turning points (R_{11} Transparency).

Prefer web-based over desktop applications Accessibility (R_{12} Accessibility) is a key factor for a visualization that is designed to be adapted by domain experts. In our case, accessibility is even more important as our target users need to share their results with their own customers. We therefore provide our visualization tool as a web application. It can be easily accessed without having to worry about installation or set-up times. A web application also lays the foundation for communicating analysis results and recommendations that goes beyond the currently used presentation slides.

Prefer objectivity over biased perception of criteria In the context of multi-attribute choice, the importance of criteria can hardly be deduced from the data themselves, as this requires the subjective judgement of the decision-maker. Each criterion is thus meaningful for the evaluation of alternatives and for the interpretation of trade-offs. In the absence of prior importance information, a visualization should make use of the same visual mapping for all criteria, unless the user explicitly requested a visual distortion.

Prefer lossless mapping over dimension reduction Dimension reduction approaches to Pareto front visualization sacrifice informativeness for the purpose of intuitive exploration and navigation. However, Dimara et al. found that *"the majority of [...] visualization tools meant to support multi-attribute choice employ lossless geometric projections"* [DBD17]. To make a multi-attribute choice, users need to be able to visually retrieve any criteria value from any alternative without interaction (R_2 Overview, R_9 Details). We therefore prefer a lossless mapping over dimension reduction. Our design target is a dozen design parameters and up to ten criteria. It is thus possible to depict all alternatives and criteria without the need for aggregation or selection of a data subset to view.

5.2. Parallel Coordinates View

PAVED's primary view shows a parallel coordinates visualization. Though the initial prototype, the scatterplot matrix, provided a lossless mapping of the Pareto front, it depicted multi-criteria trade-offs only via glyph overlay and pairwise trade-offs otherwise. However, engineers need to view alternatives as a whole (R_2 Overview). In addition, the glyph-based scatterplots did not meet the objectivity requirement, because the two criteria mapped to position are considered most important. To meet both requirements, we decided for parallel coordinates [ID90] as our final visual encoding.

Parallel coordinates present a compact and lossless two-dimensional visual representation for multi-dimensional alternatives. Different axis layouts have been proposed, e.g. many-to-many, force-directed, and three-dimensional layouts [JF15]. However, as parallel coordinates were unknown to the engineers, we stuck to the standard two-dimensional layout. For the same reason, we decided in favor of the more common vertically laid out axes (R_1 Validation). To maintain an unbiased perception of the criteria, we chose to neither invert nor scale the axes like suggested by some works [AA01]. Instead, we mark the desired direction of change by a triangular indicator at the respective end of the axis [PSTW*16].

Motivated by the need to scale with hundreds of depicted alternatives, our focus was on enhancing the perception of trade-offs and individual alternatives. For this purpose, we took advantage of standard visual encoding and interaction techniques known from literature. These techniques modify either the polylines or the axes.

Parallel coordinates are well-known for being sensitive to visual clutter, which might hide patterns and alternatives of interest. Techniques that render aggregates only are not an option, because they violate our design choice of a lossless projection. Instead, we render each individual polyline with a constant *line transparency*. Still, lossless projections are not scalable beyond a certain point. With a large number of polylines being depicted, two or more lines might intersect an axis at nearly the same position. In such a case, it is difficult to trace the individual lines. To resolve ambiguities, the user can activate *curve smoothing*, which replaces the polylines by cubic splines that interpolate the original values at the axes [GK03]. Finally, each alternative is associated with a motor topology as categorical metadata. An effective technique to support the perception of nominal data is *color-coding* [HW13]. Categorical dimensions like the motor topology can also be considered for filtering.

Gaining an overview of a Pareto front also benefits from observing the relationships between criteria (R_3 Criteria Relations). The axis order affects the pairwise relations between adjacent axes that are revealed by parallel coordinates. As *axis ordering* is a complex research problem on its own, we arrange axes by dimension order and enable users to explicitly reorder the axes according to their needs. We implemented an animated translation guided by a drag-and-drop operation, where a uniform axis spacing is reconstructed

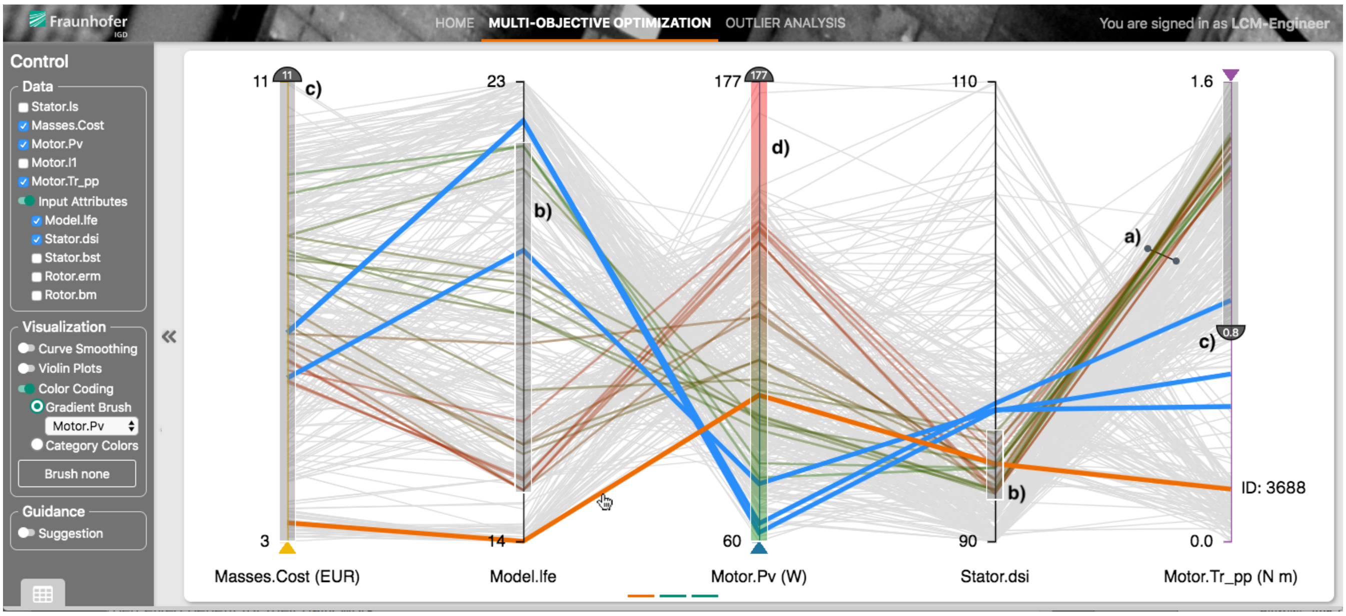

Figure 4: *This image shows different interaction modes that drive an optimization analysis. Standard brushes include a line brush for open exploration (a) and range brushes at the parameter axes (b). For criterion axes, we propose preference brushes (c), which are locked at the desired end of the axis and dragged via the labeled handle at their other end. The gradient color brush (d) is applied to* Motor.Pv*, revealing its correlation with both* Masses.Cost *and* Model.lfe*. Favorite alternatives (blue) can be stored independently of the brushes. At any time, they represent the current search result, which might also be presented to the customer. The preferred choice is the hovered alternative (orange).*

after releasing an axis [HW13]. To adjust the complexity of the parallel coordinates depending on the decision stage, the *axis visibility* can be controlled individually for design parameter and criteria axes (R_{10} Sensitivity). Design parameters are hidden by default, as large parts of an engineer's decision focus on the criteria.

5.3. Interactive Selection

Interaction is essential for an effective use of parallel coordinates. Selecting a subset of favorite alternatives for detailed analysis (R_4 Trade-offs) is complemented with filtering alternatives according to performance constraints and preferences (R_5 Filter). From a technical point of view, this corresponds to a selection by elements versus a selection by dimension values. Available interactions are indicated by a transformation of the mouse cursor. We provide a video in the supplementary material, which shows interactions by our domain expert, highlighting how practitioners use our tool.

Selecting Favorite Alternatives

To support the user in scanning through the alternatives, we provide *hovering* as the most basic interaction (Figure 4, orange). From the label to the right end of the hovered polyline, users can retrieve the alternative ID. This allows them to join insights about the alternative in focus with e.g. offline data. To select a group of alternatives, we provide a *line brush* (Figure 4a), which makes it easy to isolate alternatives with particular characteristics (R_6 User Interaction).

Users can *flag* alternatives by clicking on a polyline. Flagged options are permanently visible, even when they are not part of any other selection (Figure 4, blue). This enables a direct comparison with respect to each of their dimension values (R_8 Comparison). They can also be considered the current result of the exploration. While exploring the Pareto front, the engineers flagged alternatives to cache a small number of favorites for later in-depth comparison (R_7 Provenance). This set of superior compromises can also be kept in case the customer questions the first choice (R_{11} Transparency).

Eliminating Undesired Alternatives

We provide filtering of alternatives in the form of range brushes applied to criteria and design parameters axes (Figure 5a). As a choice in engineering design needs to satisfy multiple constraints and preferences, several of these brushes are combined to a composite brush using the logical AND operation (Figure 4b) .

On a criterion axis, the desired direction of change is known, i.e. whether low or high values are preferred. In any case, it does not make sense to exclude alternatives that are located on the desired end of an axis. We deliberately limit the interaction on the range brush in this regard to match the set of filtering operations actually needed for optimization (R_6 User Interaction). We propose the *preference brush*, a range brush that is locked at the high-quality end of the respective axis (Figure 5b). Only the low-quality end of the brush can be dragged for filtering (Figure 4c). Regarding the observed criterion, this ensures that a preference brush always includes the best solutions, while the interaction complexity is significantly reduced. A value label at the draggable end of the brush makes the current constraint settings readable for the engineer.

Each criterion axis is equipped with a preference brush, which initially covers the full axis range. Its expressiveness can be en-

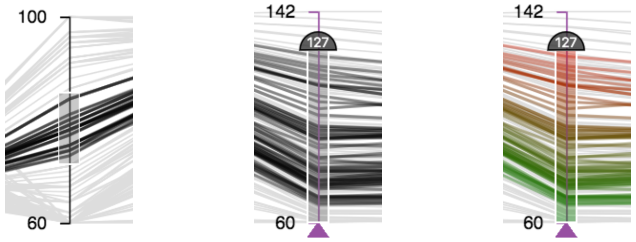

(a) *Standard Range Brush* **(b)** *Preference Brush* **(c)** *Gradient Color Brush*

Figure 5: *Parameters are filtered using a range brush (a). For criteria, brushes are locked to the high-quality end of the axis indicated by a triangle (b) and can be augmented with a color gradient (c).*

hanced by applying a red-green color gradient to its range that triggers a corresponding color-coding of the brushed polylines (Figure 5c). The *gradient color brush* was first introduced by Matkovic et al. [MJJ*05]. They mapped red to the lower and green to the upper end of the brush range to explore the influence of parameter changes on the output of an investigated system. We introduce slight modifications to help users explore where trade-offs between criteria need to be made by observing how value changes in one criterion manifest in the remaining criteria (R_3 Criteria Relations). First, we adjust the color scale such that green encodes desired and red encodes undesired criteria values. Second, the start and end colors are not assigned to the ends of the brush range, but to the ends of the axis. The meaning of colors then does not change when the brush is modified. With these modifications, desired and undesired values of a criterion can be traced across axes more easily, allowing users to observe one-to-many trade-offs (Figure 4d).

5.4. Implementation

The prototype presented in Figure 4 is a single-page web application written in *TypeScript* using the JavaScript framework *React*. The parallel coordinates view is based on the visualization library *D3.js* [BOH11]. Data are managed on the client and can be read from a JSON file or an external server. The data volumes provided by our experts can be processed with interactive response on average hardware, involving only a few seconds of initial data fetching.

6. Evaluation

The goal of this evaluation is to validate the domain usefulness of the proposed visualization in terms of effectiveness and problem-solving characteristics for experts doing their own work. By dealing with decision-making, we address a high-level cognitive task, which is difficult to measure objectively and quantitatively [TM05]. As realism in tasks, data, and users is important, we performed a qualitative field study. This study combined qualitative coding of user feedback with a quantitative usability scale. The results suggest that the tool supports the identified analysis tasks for making a multi-attribute choice from simulated design options. They also provide indications where there is potential for improvement.

6.1. Methodology

The field study was performed with motor engineers in applied research using real-world data from one of their design optimization use cases. We wanted to observe how the target users interact with the deployed visualization in their own working environment to see whether the tool met their needs. The study was conducted in the form of one-hour think-aloud sessions with one observer who was also taking notes. Each session involved a prescribed walk-through of the tool, open-ended questions about its usage, and a usability questionnaire. Five experts other than our primary domain expert participated in the study: the major contributor of the tool currently used by the engineers, a simulation expert and three experienced motor designers. A rapport was established during a preceding introduction in visual analysis where all participants were present.

The notes taken during the think-aloud sessions were analyzed using a qualitative coding methodology [SC90]. Repeated statements, ideas, or topics in the collected feedback and observations were labeled with codes extracted from the data. These codes were then grouped into more abstract categories to summarize the results of the think-aloud sessions. The categorization is aligned to a set of questions that was proposed by Lam et al. for evaluating user experience [LBI*11]. In addition to that, we quantitatively assessed the usability of our tool using the System Usability Scale (SUS), which is composed of ten statements that are rated on a Likert scale [Sau11]. On top of the five experts who already took part in our field study, we acquired a second group of three experts from another engineering field. The qualitative coding scheme together with the quantitative usability scores convey a comprehensive picture of our tool's deployment readiness level.

6.2. Results

After coding and sorting the participants' comments and our observations, we ended up with five categories: usability, useful features, missing features, limitations, and the perceived potential of visualization. The usability category includes codes that indicate the understandibility and learnability of the tool. The dynamic brushing (preference brush as well as gradient color brush) and flagging of interesting solutions were highlighted as particularly useful. This feedback was supplemented by feature suggestions like details-on-demand, automatic warnings about critical brushes, or documentation support. Hot topics regarding the potential of visualization were decision-making transparency and revisiting decisions with customers. Talking to the domain experts also revealed features that were irrelevant to them. They indicate that we over-prioritized the underlying task abstractions in the domain characterization stage. We provide more detailed reflections in Section 7.1 and 7.2.

The qualitative feedback of the target users also uncovered a comprehensibility issue. Three of five participants were confused by the brushes being locked to one end of the axes. Most of the selection rectangles that they encounter in their daily or working life can be modified with respect to all directions. However, explaining the reasoning behind the proposed preference brushes led the participants to reconsider their initial expectation and to confirm our underlying abstraction: *"you're right, I cannot think of any situation, where I would want to move the other side, too"* (A5).

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Total
A1	7.5	10	10	7.5	10	10	7.5	10	7.5	10	90
A2	10	10	10	10	7.5	10	10	10	7.5	10	95
A3	10	10	10	7.5	7.5	10	10	10	10	7.5	92.5
A4	7.5	7.5	10	10	7.5	7.5	7.5	10	7.5	10	85
A5	10	7.5	10	7.5	7.5	7.5	10	10	10	10	90
B1	5	10	10	10	7.5	5	10	10	10	10	87.5
B2	5	10	10	10	7.5	7.5	10	10	7.5	10	87.5
B3	7.5	10	7.5	7.5	10	7.5	10	10	7.5	10	87.5
Avg	7.8	9.4	9.7	8.8	8.1	8.1	9.4	10	8.4	9.7	**89.4**

Table 1: *Results of the System Usability Scale [Sau11] with two groups A and B of domain experts. The tool achieved a total score of 89.4 out of 100. Interest in frequent use (Q1) received the lowest score, while ease of use (Q8) was rated particularly high.*

According to the experts, the scatterplots in their current tool are well-suited for observing the progress of genetic optimization and steering it. However, they like our visualization tool for making the actual preferential choice on the resulting Pareto front. They specify different reasons for this. Two of them stated that our tool provides them with a more intuitive and flexible brushing functionality. Another one preferred scatterplots for pairwise trade-offs, because they convey how a Pareto front is bent. Still, he agreed that higher-dimensional trade-offs require techniques like parallel coordinates. When comparing both tools, the experts became particularly aware of how different visualizations support different kinds of tasks. Consequently, one of them suggested to combine the strengths of both, i.e. to simultaneously observe scatterplots and parallel coordinates without having to switch views.

The quantitative results of the usability survey suggest that our exploration tool provides an *excellent* usability, according to the adjective equivalent of the achieved SUS score [BKM09]. With a score of 89.4, we found the usability of our tool to be highly above average, which is reflected by a score of 68 out of 100 [Sau11]. We present the individual scores broken down by question in Table 1. We noticed that our tool scored highest on ease of use (Q8), where all participants agreed on the strongest possible approval. In contrast, the statement about the participants' prognosis of using the tool frequently (Q1) received the lowest ratings (which were still agreeing in total). Here, the two groups' ratings differed significantly (9 versus 5.8 out of 10). A possible explanation might be the different visualization and domain background of either group, which affects the tool's perceived benefit for their daily work.

7. Reflections

As Sedlmair et al. point out, contributions that make design studies useful for other visualization researchers focus on various aspects of the problem domain, a validated visualization tool, or reflections on design guidelines [Sed16]. Meyer and Dykes particularly stress that the knowledge acquired through a design study is highly subjective and needs to be viewed in the context of its generation [MD19]. Inspired by their proposed subdivision of contributions into three topics, we provide our reflections on 1) the problem domain, 2) visualization idioms, and 3) methodological guidelines.

7.1. Problem Domain

While being strongly interested in the potential of visualization, design engineers are rarely visualization experts themselves. Research on multi-criteria optimization in engineering design mainly focuses on advancing simulation and optimization algorithms. As the ability to simulate even larger design spaces increases, the importance of visualization to make sense out of the growing Pareto fronts raises to the same extent. In line with the insight "simpler dashboards are better" of Arbesser et al. [AMKP17], we realized that our simple interactive visualization provided a clear benefit for the domain experts, although these techniques are considered a standard in our domain. We thus argue for a greater consideration of well-known visualizations like parallel coordinates under careful consideration of their practicability in new industry applications.

We found the most important design requirements to be 1) the ability to retrieve any criteria value of any design from the visual encoding (R_2 Overview) and 2) the transparency of the decision-making process (R_{11} Transparency). The latter is important because a comprehensible decision process *"helps to prove plausibility and justify the decision"* (A3). This is closely related to the need to communicate and collaboratively revisit a decision. Design engineers need to be able to explain how they reached their design decision because *"an understanding of the optimization problem and selection process is highly important to customers"* (A4). From their experience, an interactive exploration of alternatives is highly beneficial for such explanations. The interactive visualization even offers the potential to involve customers in the decision-making (A3).

The engineers did not consider guidance relevant for an exploration. We suggested an automatic highlighting of the next best solution to guide them towards interesting regions of the solution space. However, they prefer to explore the available alternatives by themselves using the brushing mechanism. This might originate from their awareness that the simulation model itself can contain inaccuracies that a fully automated optimization would overlook.

Our collaboration and discussions made the domain experts think differently about their tasks and workflow. Interacting with the dynamic brush led them to realize why they perceived the interaction in their current tool as not very straightforward: because it involved quite a few mouse clicks and its effects were thus not instantly visible. The brushing mechanism also made the simulation expert think about applying fuzzy logic to the brushes: *"maybe vague preferences are better represented by a fuzzy selection"* (A3). In our domain, this is known as smooth brushing [DH02]. We were also pointed towards optimization scenarios involving interdependent components. They remain an open challenge for now, thus offering a promising perspective for continued collaboration.

We have also been approached by a party from the rail supply industry, who is interested in using our tool for optimizing the production process of transportation pallets. Their interest confirms that our abstraction seems to be at the right level, because our visualization can adapt to different optimization problems from different domains, targeting both products and processes. However, this is subject to a formal evaluation. Also, parallel coordinates rarely scale well with the number of data items. This issue is partly mitigated by the fact that the number of items in focus is reduced very early during exploration via filtering.

7.2. Visualization Idioms

Motivated by the recommendation that *"studies [of parallel coordinates] in new application areas should be encouraged"* [JF15], we discuss aspects of our visual encoding, interaction techniques, and envisioned analysis work flow regarding their acceptance by design engineers. We also comment on visualization design guidelines.

Johansson and Forsell have found parallel coordinates to be *"advantageous to state-of-the-art techniques when introduced in a new application area"* [JF15]. Our findings align with this. The domain experts quickly became familiar with the visual encoding. Due to the lossless projection, they had no difficulties in gaining an overview of the available multi-criteria alternatives. They particularly appreciated the brushing mechanism and observing its direct effect on the selection of alternatives. This even seemed to have outweighed the well-known issue of parallel coordinates being sensitive to visual clutter. For optimization, we would like to promote the preference brush as a simplification of composite brushes.

The existing analysis tool for the design and optimization of motors is built around an interactive scatterplot matrix depicting pairwise trade-offs. Scatterplots are known to convey correlations more effectively than parallel coordinates [LMVW10]. Our experts also stated that they prefer to observe pairwise trade-offs in a scatterplot. Still, due to their ability to convey an overview, parallel coordinates were rated high as a complement to the traditional scatterplots. This aligns with Yuan's et al. combination of scatterplots and parallel coordinates to exploit the strengths of both [YGX*09]. In their recent study, Dimara et al. found that tabular layouts were preferred over parallel coordinates for decision-making tasks [DBD17]. However, tabular layouts often require users to explicitly express their criteria preferences for ranking purposes [GLG*13]. Still, some experts confirmed the relevance of tabular visualizations: they would have appreciated a linked brushing functionality for the table view.

Some visual encodings were not effective in this domain. Although the radial bar charts provide a compact representation of individual alternatives, the engineers were not satisfied with the overview of the criteria space. The curve smoothing that should support users in tracing lines that intersect the axes in common points was not considered relevant for the perception of alternatives. The interactive translation of axes was initially discussed controversially, but some engineers quickly adapted to using this feature. For one-to-many correlations the experts commonly appreciated the gradient color brush as *"intuitive"* (A4) and *"practical"* (A5). In particular, the experts used this brush to observe how changing values in one criterion affected the remaining criteria.

7.3. Methodological Guidelines

Real data being available from the very beginning of the project helped a lot in developing an understanding of the problem statement and identifying valid abstractions that shaped the design of our tool early in the process. Our domain experts committed a lot of time for problem analysis and design discussions. This commitment was in large parts based on an exceptional intrinsic motivation that stemmed from their personal interest in visualization as well as enjoying problem and design discussions through a positive rapport between researchers and domain experts.

Apart from that, our collaboration was effective for two more reasons. First, both parties were willing to familiarize themselves with the subjects of the other party. The domain experts already knew basic visualization and interaction concepts, which significantly reduced the initial knowledge gap. Second, we encouraged meetings in person for joint sketching on the same whiteboard or piece of paper to generate and evaluate ideas. This led to results more efficiently than having one party prepare content that is reviewed by the other. By constantly providing prototypes, we were able to reinforce the experts' engagement and keep their attention. In the end, our discussions with the domain experts were so inspiring that we even identified an entirely new problem that poses an interesting research question in both domains.

We should have listened more carefully when the domain experts encouraged the use of parallel coordinates shortly after our collaboration started. We initially missed this suggestion due to a blind spot that grew from our assumption that building upon familiar visualizations would 1) avoid the pitfall of ignoring practices that work well and 2) keep the tool easy to learn. However, from this detour, we had to acknowledge that, despite their limited visualization background, our domain experts had a meaningful understanding of their visualization needs. Consequently, when performing user-centered design, we learned to consider the users' suggestions more strongly, even if we might feel our expertise being underrated.

8. Conclusion

In this work, we present the results of a design study on Pareto front visualization in the field of engineering design. In close collaboration with electromechanical engineers, we studied their tasks and needs related to making a multi-attribute choice. The resulting problem characterization guided our development of PAVED, an interactive parallel coordinates visualization that enables the exploration of design alternatives, which are characterized by about a dozen design parameters and up to ten criteria. The visualization supports engineers in applying both formal constraints and informal preferences as they learn what level of performance is achievable under different conditions. This allows them to understand the trade-offs involved, which is essential for justifying the final choice to their customers. Although designed for engineers, our visualization can also be used by other decision-makers like consumers or professional buyers, policy makers, or event managers. A qualitative field study indicates the usefulness of our approach for performing real-world optimization, which also manifests in the results of a usability testing. Finally, we reflect on domain characteristics, visualization design, and methodological considerations and showed the benefit of a well-known visualization like parallel coordinates for industry applications. Our future research will focus on adapting the parallel coordinates visualization to support the optimization of systems consisting of interdependent components.

9. Acknowledgements

This work was funded by the EC H2020 programme under grant agreement n°768892. This work has been supported by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian federal government and the federal state of Upper Austria. The dataset presented is courtesy of LCM.

References

[AA01] ANDRIENKO G., ANDRIENKO N.: Constructing parallel coordinates plot for problem solving. In *1st International Symposium on Smart Graphics* (2001), pp. 9–14. 4, 6

[AMKP17] ARBESSER C., MÜHLBACHER T., KOMORNYIK S., PIRINGER H.: Visual analytics for domain experts: Challenges and lessons learned. In *Proceedings of the Second International Symposium on Virtual Reality & Visual Computing* (2017). 4, 9

[BC03] BAGAJEWICZ M., CABRERA E.: Pareto optimal solutions visualization techniques for multiobjective design and upgrade of instrumentation networks. *Industrial & Engineering Chemistry Research 42*, 21 (2003), 5195–5203. 4

[BHGK14] BEHAM M., HERZNER W., GRÖLLER M. E., KEHRER J.: Cupid: Cluster-based exploration of geometry generators with parallel coordinates and radial trees. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1693–1702. 4

[BKM09] BANGOR A., KORTUM P., MILLER J.: Determining what individual sus scores mean: Adding an adjective rating scale. *Journal of Usability Studies 4*, 3 (2009), 114–123. 9

[BMPM12] BOOSHEHRIAN M., MÖLLER T., PETERMAN R. M., MUNZNER T.: Vismon: Facilitating analysis of trade-offs, uncertainty, and sensitivity in fisheries management decision making. In *Computer Graphics Forum* (2012), vol. 31, pp. 1235–1244. 4

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D^3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2301–2309. 8

[BPFG11] BERGER W., PIRINGER H., FILZMOSER P., GRÖLLER E.: Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. In *Computer Graphics Forum* (2011), vol. 30, pp. 911–920. 4

[BQP*15] BASOLE R. C., QAMAR A., PARK H., PAREDIS C. J. J., MCGINNIS L. F.: Visual analytics for early-phase complex engineered system design support. *IEEE Computer Graphics and Applications 35*, 2 (2015), 41–51. 4

[BSM*13] BERGNER S., SEDLMAIR M., MÖLLER T., ABDOLYOUSEFI S. N., SAAD A.: Paraglide: Interactive parameter space partitioning for computer simulations. *IEEE Transactions on Visualization and Computer Graphics 19*, 9 (2013), 1499–1512. 4

[CAS*13] CHEN S., AMID D., SHIR O. M., LIMONAD L., BOAZ D., ANABY-TAVOR A., SCHRECK T.: Self-organizing maps for multi-objective Pareto frontiers. In *IEEE Pacific Visualization Symposium (PacificVis)* (2013), pp. 153–160. 4, 5

[DBD17] DIMARA E., BEZERIANOS A., DRAGICEVIC P.: Conceptual and methodological issues in evaluating multidimensional visualizations for decision support. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 749–759. 3, 4, 6, 10

[DH02] DOLEISCH H., HAUSER H.: Smooth brushing for focus+context visualization of simulation data in 3d. In *Proceedings of the 10th International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media (WSCG)* (2002), pp. 147–154. 9

[DI13] DUAN Y., IONEL D. M.: A review of recent developments in electrical machine design optimization methods with a permanent-magnet synchronous motor benchmark study. *IEEE Transactions on Industry Applications 49*, 3 (2013), 1268–1275. 2

[FPL05] FLEMING P. J., PURSHOUSE R. C., LYGOE R. J.: Many-objective optimization: An engineering design perspective. In *International Conference on Evolutionary Multi-Criterion Optimization* (2005), Springer, pp. 14–32. 4

[GBS*99] GOEL A., BAKER C., SHAFFER C. A., GROSSMAN B., HAFTKA R. T., MASON W. H., WATSON L. T.: Vizcraft: A multi-dimensional visualization tool for aircraft configuration design. In *Proceedings of the conference on Visualization '99* (1999), pp. 425–428. 4

[GK03] GRAHAM M., KENNEDY J.: Using curves to enhance parallel coordinate visualisations. In *Proceedings on 7th International Conference on Information Visualization (IV)* (2003), IEEE, pp. 10–16. 6

[GLG*13] GRATZL S., LEX A., GEHLENBORG N., PFISTER H., STREIT M.: Lineup: Visual analysis of multi-attribute rankings. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2277–2286. 4, 10

[Hor96] HORN J.: *Multicriteria Decision Making and Evolutionary Computation*. Tech. rep., 1996. 1

[HW13] HEINRICH J., WEISKOPF D.: State of the art of parallel coordinates. In *Eurographics (STARs)* (2013), pp. 95–116. 6, 7

[ID90] INSELBERG A., DIMSDALE B.: Parallel coordinates: A tool for visualizing multi-dimensional geometry. In *Proceedings of the 1st Conference on Visualization '90* (1990), IEEE Computer Society Press, pp. 361–378. 6

[JF15] JOHANSSON J., FORSELL C.: Evaluation of parallel coordinates: Overview, categorization and guidelines for future research. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2015), 579–588. 6, 10

[KIPS13] KIPOUROS T., INSELBERG A., PARKS G., SAVILL A. M.: Parallel coordinates in computational engineering design. In *54th Structures, Structural Dynamics, and Materials Conference* (2013), p. 1750. 4

[Koh01] KOHONEN T.: *Self-Organizing Maps*. Springer, 2001. 4

[KW08] KORHONEN P., WALLENIUS J.: Visualization in the multiple objective decision-making framework. In *Multiobjective Optimization*. Springer, 2008, pp. 195–212. 4

[LBI*11] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics 18*, 9 (2011), 1520–1536. 8

[LBK13] LOTOV A. V., BUSHENKOV V. A., KAMENEV G. K.: *Interactive Decision Maps: Approximation and Visualization of Pareto Frontier*, vol. 89. Springer Science & Business Media, 2013. 4

[LM08] LOTOV A. V., MIETTINEN K.: Visualizing the Pareto frontier. In *Multiobjective Optimization*. Springer, 2008, pp. 213–243. 4

[LMVW10] LI J., MARTENS J.-B., VAN WIJK J. J.: Judging correlation from scatterplots and parallel coordinate plots. *Information Visualization 9*, 1 (2010), 13–30. 10

[MD19] MEYER M., DYKES J.: Criteria for rigor in visualization design study. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 87–97. 9

[MGH*07] MARSAW A., GERMAN B., HOLLINGSWORTH P., MAVRIS D., CSONKA S.: An interactive visualization environment for decision making in aircraft engine preliminary design. In *45th AIAA Aerospace Sciences Meeting and Exhibit* (2007), p. 1333. 4

[MJJ*05] MATKOVIC K., JELOVIC M., JURIC J., KONYHA Z., GRACANIN D.: Interactive visual analysis and exploration of injection systems simulations. In *Proceedings of the IEEE Visualization 2005* (2005). 4, 8

[MLMP17] MÜHLBACHER T., LINHARDT L., MÖLLER T., PIRINGER H.: Treepod: Sensitivity-aware selection of Pareto-optimal decision trees. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 174–183. 4

[Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009), 921–928. 1, 5

[OJC05] OBAYASHI S., JEONG S., CHIBA K.: Multi-objective design exploration for aerodynamic configurations. In *35th AIAA Fluid Dynamics Conference and Exhibit* (2005), p. 4666. 4

[PBCR11] PRETORIUS A. J., BRAY M.-A., CARPENTER A. E., RUDDLE R. A.: Visualization of parameter space for image analysis. *IEEE*

Transactions on Visualization and Computer Graphics 17, 12 (2011), 2402–2411. 4

[PPBJ93] PAYNE J. W., PAYNE J. W., BETTMAN J. R., JOHNSON E. J.: *The Adaptive Decision Maker*. Cambridge University Press, 1993. 3

[PSTW*16] PAJER S., STREIT M., TORSNEY-WEIR T., SPECHTENHAUSER F., MÖLLER T., PIRINGER H.: Weightlifter: Visual weight space exploration for multi-criteria decision making. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2016), 611–620. 4, 6

[PWB*09] POTTER K., WILSON A., BREMER P.-T., WILLIAMS D., DOUTRIAUX C., PASCUCCI V., JOHNSON C. R.: Ensemble-vis: A framework for the statistical visualization of ensemble data. In *IEEE International Conference on Data Mining Workshops* (2009), pp. 233–240. 4

[Sau11] SAURO J.: Measuring usability with the system usability scale (sus). `http://measuringu.com/sus`, 2011. Accessed 05-Nov-2019. 8, 9

[SBD*16] SILBER S., BRAMERDORFER G., DORNINGER A., FOHLER A., GERSTMAYR J., KOPPELSTÄTTER W., REISCHL D., WEIDENHOLZER G., WEITZHOFER S.: Coupled optimization in MagOpt. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering 230*, 4 (2016), 291–299. 2

[SC90] STRAUSS A. C., CORBIN J.: Basics of qualitative research: Grounded theory procedures and techniques, 1990. 8

[Sed16] SEDLMAIR M.: Design study contributions come in different guises: Seven guiding scenarios. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), ACM, pp. 152–161. 9

[SFMB12] SEDLMAIR M., FRANK A., MUNZNER T., BUTZ A.: Relex: Visualization for actively changing overlay network specifications. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2729–2738. 4

[SHB*14] SEDLMAIR M., HEINZL C., BRUCKNER S., PIRINGER H., MÖLLER T.: Visual parameter space analysis: A conceptual framework. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2161–2170. 4

[SKW98] SHAFFER C. A., KNILL D. L., WATSON L. T.: Visualization for multiparameter aircraft designs. In *Proceedings Visualization'98* (1998), IEEE, pp. 491–494. 4

[SKW*18] SILBER S., KOPPELSTÄTTER W., WEIDENHOLZER G., SEGON G., BRAMERDORFER G.: Reducing development time of electric machines with symspace. In *8th International Electric Drives Production Conference (EDPC)* (2018), IEEE, pp. 1–5. 2

[SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2431–2440. 1

[SOL*15] SORGER J., ORTNER T., LUKSCH C., SCHWÄRZLER M., GRÖLLER E., PIRINGER H.: Litevis: Integrated visualization for simulation-based decision support in lighting design. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2015), 290–299. 4

[STDS95] SPENCE B., TWEEDIE L., DAWKES H., SU H.: Visualisation for functional design. In *Proceedings of Visualization 1995 Conference* (1995), IEEE, pp. 4–10. 3

[TM05] TORY M., MÖLLER T.: Evaluating visualizations: Do expert reviews work? *IEEE Computer Graphics and Applications 25*, 5 (2005), 8–11. 8

[TM12] TANAHASHI Y., MA K.-L.: Design considerations for optimizing storyline visualizations. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2679–2688. 5

[TWSM*11] TORSNEY-WEIR T., SAAD A., MOLLER T., HEGE H.-C., WEBER B., VERBAVATZ J.-M., BERGNER S.: Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 1892–1901. 4

[TWSM15] TORSNEY-WEIR T., SEDLMAIR M., MÖLLER T.: Decision making in uncertainty visualization. In *IEEE VIS Workshop on Visualization for Decision Making under Uncertainty (VDMU)* (2015). 3

[Ull01] ULLMAN D. G.: Robust decision-making for engineering design. *Journal of Engineering Design 12*, 1 (2001), 3–13. 3

[YGX*09] YUAN X., GUO P., XIAO H., ZHOU H., QU H.: Scattering points in parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009), 1001–1008. 10

[ZBL*13] ZĂVOIANU A.-C., BRAMERDORFER G., LUGHOFER E., SILBER S., AMRHEIN W., KLEMENT E. P.: A hybrid soft computing approach for optimizing design parameters of electrical drives. In *Soft Computing Models in Industrial and Environmental Applications*. Springer, 2013, pp. 347–358. 2

[ZWC*17] ZHAO X., WU Y., CUI W., DU X., CHEN Y., WANG Y., LEE D. L., QU H.: Skylens: Visual analysis of skyline on multi-dimensional data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 246–255. 4