VISUAL ANALYSIS FOR MULTI-ATTRIBUTE CHOICE

Vom Fachbereich Informatik
der Technischen Universität Darmstadt
genehmigte

DISSERTATION

zur Erlangung des akademischen Grades
Doktor-Ingenieurin (Dr.-Ing.$^{\text{in}}$)

vorgelegt von

M.SC. LENA CIBULSKI

Referenten der Arbeit:

Prof. Dr.-Ing. Jörn Kohlhammer
Technische Universität Darmstadt

Prof. Dr. techn. Dr.-Ing. eh. Dieter W. Fellner
Technische Universität Darmstadt

Univ.-Prof. Dipl.-Ing. Dr. techn. Eduard Gröller
Technische Universität Wien

Tag der Einreichung: 09.02.2024
Tag der Prüfung: 25.03.2024

D17 - Darmstadt 2024

Lena Cibulski
Visual Analysis for Multi-Attribute Choice
Dissertation, Darmstadt, Technische Universität Darmstadt

Jahr der Veröffentlichung: 2024
Tag der mündlichen Prüfung: 25.03.2024
URN: urn:nbn:de:tuda-tuprints-277869
URL: https://tuprints.ulb.tu-darmstadt.de/id/eprint/27786

If we knew what it was we were doing,
it would not be called *research*, would it?

— Albert Einstein

If you can look into the seeds of time,
and say which grain will grow and which will not,
speak then to me.

— William Shakespeare, Macbeth

We are tempted to ask how all the effort
we invest in research will shape the future.

Foretelling the future of research efforts is
like observing a number of seeds and predicting
which will take hold and sprout, and which will not.

But we cannot know which seeds of time will grow.
We can only work hopefully, doing our best
to push the boundaries of what we know
that our contributions may become seeds
that grow into lasting benefits.

ABSTRACT

This dissertation presents findings from problem-driven research that centers around the design of visualization tools to assist experts in making data-informed choices.

Identifying the most preferred solution among many alternatives is a common task in our everyday and professional lives. Pivotal information is usually hidden in the data and visualization research has long treated decision-making as a data comprehension task. To arrive at a decision, however, the understanding of patterns in the data needs to be synthesized with subjective judgments. Existing visualization tools do not target this synthesis and many approaches focus on simplified decisions tasks. As a result, their relevance and applicability in real-world settings might be limited.

This dissertation promotes field research to investigate the cognitive processes underlying real-world decisions and to operationalize them for the design and validation of decision support tools. Being based on a close collaboration with real decision-makers, it provides an emphasis on decision processes that problem-driven visualization research did not have before. By synthesizing the collected real-world experience with concepts from human science, it also contributes to making decision models and theories usable for visualization design.

This dissertation refines the existing multi-attribute choice definition by describing it as a constructive problem where preferences are incrementally formed at the actual time of choice. It further proposes a characterization scheme to help visualization researchers concretize the decision problem to design for. Finally, going into the field revealed a novel type of constructive decision problem, which this dissertation defines as co-dependent choices. As theoretical contributions, these formalizations make the design space of decision tasks more tangible. A knowledge elicitation method is adapted from cognitive science to systematically detail the knowledge, experience, and cognitive tasks underlying current decision-making practices. As a methodological contribution, this introduces a decision-oriented way of conducting problem characterizations.

As technical contributions, this dissertation presents two design studies. Their results demonstrate the relevance and applicability of the proposed concepts within and beyond the studied decision contexts. PAVED provides a simple yet effective means for decision-makers to construct and apply preferences as they learn what level of performance is achievable. Its extension COMPO*SED is the first tool that helps decision-makers explore the side effects of co-dependent choices. Their usefulness has been confirmed with domain experts on

their day-to-day decision problems. The long-term benefit of PAVED is indicated by the adoptions recorded after four years.

Through user-centered design, this dissertation addresses the lack of discourse on validated visualization tools that are dedicated to assist expert choices in the wild. Its theoretical, methodological, and technical contributions shape the understanding of decision-related activities on large data sets and how to support them with visualization. As task clarity, design guidelines, and real-world experience with decision support tools evolve, more rigorous claims regarding decision-making as a core goal of visualization research will be possible. The research presented in this dissertation is an important step in this direction.

ZUSAMMENFASSUNG

Die vorliegende Arbeit stellt die Ergebnisse anwendungsorientierter Forschung dar, in deren Mittelpunkt die Entwicklung von Visualisierungswerkzeugen steht, die Expertinnen und Experten bei datenbasierten Entscheidungen unterstützen.

In unserem alltäglichen und beruflichen Leben müssen wir häufig eine bevorzugte Option aus einer Menge von Alternativen auswählen. Dabei ist die entscheidende Information oft in einem Datensatz verborgen. Lange Zeit spielte für die Visualisierungsforschung das Datenverständnis die zentrale Rolle bei der Entscheidungsfindung. Um jedoch zu einer Entscheidung zu gelangen, muss das Verständnis von Mustern in den Daten mit subjektiven Einschätzungen zusammengeführt werden. Existierende Visualisierungswerkzeuge zielen nicht auf dieses Zusammenspiel ab und viele Ansätze konzentrieren sich auf vereinfachte Entscheidungsaufgaben. Ihre Relevanz und Anwendbarkeit unter realen Bedingungen sind daher womöglich begrenzt.

Diese Arbeit nutzt die Methodik der Feldforschung, um die zugrunde liegenden kognitiven Prozesse von Entscheidungen zu untersuchen und sie für die Entwicklung von Werkzeugen zur Entscheidungsunterstützung zugänglich zu machen. Durch eine enge Zusammenarbeit mit Entscheidungsträgerinnen und Entscheidungsträgern in der Praxis und durch den besonderen Fokus auf Entscheidungsprozesse geht die Arbeit über den bisherigen Stand der anwendungsorientierten Visualisierungsforschung hinaus. Indem die gesammelte Praxiserfahrung mit Konzepten aus der Humanwissenschaft zusammengeführt wird, trägt die Arbeit außerdem dazu bei, Entscheidungsmodelle und -theorien für die Entwicklung von Visualisierungen nutzbar zu machen.

Die bestehende Definition von Multi-Attribute Choices wird in dieser Arbeit um eine Einordnung als sogenanntes konstruktives Problem erweitert, bei dem Präferenzen erst zum Zeitpunkt der tatsächlichen Entscheidung formiert werden. Außerdem schlägt die Arbeit ein Charakterisierungsschema vor, mit dem Visualisierungsforscherinnen und Visualisierungsforscher ein gegebenes Entscheidungsproblem konkreter als bisher beschreiben können. Schließlich brachte die praxisnahe Forschung ein neuartiges Entscheidungsproblem hervor, das in dieser Arbeit als Co-Dependent Choices definiert wird. Als theoretische Beiträge tragen diese Formalisierungen zu einer Konkretisierung des Gestaltungsraumes von Entscheidungsaufgaben bei. Eine aus der Kognitionswissenschaft übertragene Methode zur Wissenserhebung dient dazu, das Wissen, die Erfahrung und die kognitiven Aufgaben, die aktuellen Entscheidungspraktiken zugrunde liegen, systematisch zu erfassen. Als methodologischer Beitrag eröffnet dies eine auf Ent-

scheidungen ausgerichtete Art und Weise der Charakterisierung eines gegebenen Problems.

Als technische Beiträge präsentiert diese Arbeit zwei Designstudien, deren Ergebnisse auf die Relevanz und Anwendbarkeit der hier vorgestellten Konzepte innerhalb sowie außerhalb der untersuchten Entscheidungskontexte hindeuten. PAVED bietet Entscheidungsträgerinnen und Entscheidungsträgern einen einfachen, aber effektiven, Weg, ihre Präferenzen zu formen und anzuwenden, basierend auf der Erkenntnis, welche Leistung unter welchen Bedingungen erreichbar ist. Die Erweiterung COMPO*SED ist das erste Werkzeug, das Entscheidungsträgerinnen und Entscheidungsträgern dabei hilft, die Wechselwirkungen von abhängigen Entscheidungen zu ergründen. Die Nützlichkeit beider Werkzeuge wurde von Domänenexpertinnen und Domänenexperten im Kontext ihrer alltäglichen beruflichen Entscheidungen bestätigt. Eine durch die Domänenexpertinnen und Domänenexperten selbst initiierte Nutzung in ihrem Arbeitsalltag nach vier Jahren zeigt den Langzeitnutzen von PAVED.

Mittels User-Centered Design begegnet diese Arbeit dem fehlenden Diskurs über validierte Visualisierungswerkzeuge zur gezielten Unterstützung von Expertinnen und Experten bei deren Entscheidungen. Die theoretischen, methodologischen und technischen Beiträge prägen das allgemeine Verständnis von Entscheidungsaktivitäten auf großen Datenbeständen und von deren Unterstützungsmöglichkeiten mittels Visualisierung. Mit dem Fortschreiten von Aufgabenklarheit, Gestaltungsrichtlinien und praktischen Erfahrungen mit Entscheidungsunterstützungswerkzeugen werden immer genauere Aussagen hinsichtlich der Entscheidungsunterstützung als Hauptziel der Visualisierungsforschung möglich sein. Die hier vorgestellten Forschungsarbeiten sind ein wichtiger Schritt in diese Richtung.

I've never scored a goal without receiving a pass from somebody else.

— Abby Wambach, soccer player

ACKNOWLEDGMENTS

THIS dissertation is a milestone on a journey that started with me typing the word "visual" into a search engine for study programs. Never would I have thought to be taken by visualization research that early. I am sincerely thankful for the many inspiring people on my way, who had a lasting impact on my academic and personal life.

My sincere gratitude goes to my primary advisor *Professor Jörn Kohlhammer*. Jörn, thank you for giving me the space to pursue my own ideas while providing competent advice and sensitive guidance whenever needed. My thanks also go to my secondary advisor *Professor Dieter Fellner* for constantly following up on the status of my research. I also thank *Professor Eduard Gröller*, not only for completing the advisory body but also for sharing his triangular world view with me and for welcoming me in Vienna more than once.

Most ideas grow better when they are exposed to someone else. In this sense, I would like to express my appreciation to all my colleagues at IGD, who created a positive and supportive working atmosphere. I particularly thank *Thorsten May* for ingenious comments making me laugh and learn at the same time, *Arjan Kuijper* for helping me with the glue and coating towards the end, and *Marija Dutz* for endless conversations on so many topics. My thanks also go to *Michel Krämer* for his constantly wise scientific, technical, and personal support. I also thank *Patricia Häg* for patiently handling each of my administrative concerns. Further, I would like to thank my fellow visualization researchers and co-authors: *Jürgen Bernard* for highly energetic academic conversations, *Johanna Schmidt* for being a role model, *Evanthia Dimara* for a lot of inspiration, *Tobias Schreck* for providing an appreciative atmosphere, and *Wolfgang Aigner* for taking me seriously from the beginning.

Problem-driven research is not possible without open-minded partners in an application domain. Thanks for being the best domain experts I could have wished for: *Hubert Mitterhofer*, *Edmund Marth*, *Siegfried Silber*, *Gerald Jungmayr*, *Gerd Bramerdorfer*, and all colleagues at LCM and EAL. Thank you for your support from start to end, your enthusiasm, and all the insights into your profession. I am also grateful for your welcoming hospitality during my research stay, which was supported by the K2 Center Symbiotic Mechatronics. Hubert, special thanks for trusting me when we met at the beginning of my doctorate.

I would also like to express my gratitude to *Bernhard Preim* and *Holger Theisel*. Bernhard, you have been a true mentor to me from early on. Together with Holger, you have sparked my interest in research and continued to encourage my academic growth in the visualization community. My first steps in visualization research were also accompanied by *Krešimir Matković*. Kreso, I still draw from your advice during my half-year stay in Vienna, granting me a peek behind the scenes of publications and academic research. I am very happy to have experienced my first-ever conference talk in your company.

Finally, I thank my parents *Christiane* and *Armin Cibulski* and my dear friend *Oscar*, whose company will always be a safe space for me. They have constantly managed to lure me out of my comfort zone when all I wanted to do was hide myself away. I would also like to deeply thank my partner *Franz Kuntke*, whose supportive hand I can always feel on my back, not only when pushing me uphill on my singlespeed. I am looking forward to all there is to come.

CONTENTS

Part I

PRELIMINARIES

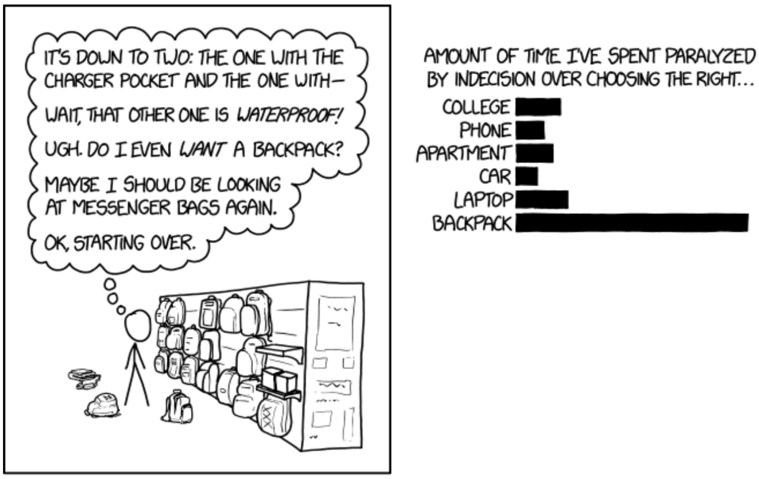

Figure 1.1: This comic nicely summarizes the difficulties with making choices, even in seemingly simple settings. Adapted from "Backpack Decisions" by Randall Munroe, licensed under CC BY-NC 2.5.

1 INTRODUCTION

We face many choices in our personal and professional lives. Some of them are trivial, like whether to have coffee or tea in the morning, others require more intellectual effort. As a tourist, we choose a hotel for our holiday stay. As a consumer, we decide which product to buy (Figure 1.1, left). As an engineer, we choose a prototype to be taken to production. As an investor, we choose stocks that hold parts of our financial resources. Some choices involve higher stakes and more careful consideration than others.

1.1 MOTIVATION

While the particular settings of these choices might differ, they have a common ground: an objectively optimal solution does not exist. Instead, options meet multiple properties differently well, such that each option is valid in its own way. To identify the best compromise for our needs and not overlook a solution, we usually compile ourselves large sets of options to choose from. Searching for accommodations in Frankfurt, Germany, on a booking platform yields 287 results[1]. Searching for coffee machines on a price comparison portal results in 526 offers[2]. Simulation in engineering generates thousands of product variations. Searching for sustainable ETFs on the stock market pro-

1 booking.com as of 30.12.2023
2 geizhals.de as of 30.12.2023

duces 570 results[3]. These options can be represented as multivariate data in a data table, where rows are options and columns are the properties associated with each option.

Choosing an option is difficult for multiple reasons. It is highly exploratory and involves trade-offs between conflicting aspects. This requires human judgment. However, human reasoning about the superiority of an option is not necessarily consistent with rationality [41, 63]. The limited cognitive capability of our human brains makes a full cost-benefit analysis unfeasible. This leaves us with no choice but to form heuristics and select a satisfactory rather than an optimal solution. Which option is chosen is inherently subjective. It depends on individual goals, preferences, and experience. It is subject to knowledge, interpretation, and intuition. These prerequisites also manifest differently in professional decision-makers, who are trained to repeatedly make expert decisions in their field, and casual decision-makers, who face personal decisions at different points in their lives. In any case, goals and preferences might even shift as we learn what level of performance is achievable under different conditions.

To arrive at a choice, we can employ different strategies. We could, for example, proceed with the first working solution, follow our internal gut feeling, or focus on selected pieces of information only[4]. While such strategies might be valid and helpful in certain contexts, the closest to a factually sound decision basis might be a systematic analysis of the available data. However, at the other side of the spectrum, an extensive consideration of data might provoke slowed or even stopped choice progress due to an overload with options and the tendency to overthink[5] (Figure 1.1, right).

Any choice we make is based on perceiving information in the data that shifts our tendency in one or the other direction. However, understanding quantitative information, even in simpler data sets, can be difficult and might require comprehension of statistical measures. In such settings that involve many options, conflicting goals, and subjective judgments, it is all the more important to present the data in a way that facilitates meaningful interpretation and exploration. Thus, to handle the challenge of making choices, we need visualization. Designing effective visualization support benefits from data science, the study of extracting patterns and meaning from data, but also from cognitive science, the study of human beings and how they make decisions. Data visualization can build upon both disciplines to provide human choice-makers with a means to consume information effectively without the need for extensive training.

---

3 justetf.com as of 30.12.2023

4 Making decisions without systematic study and reflection is called *extinction by instinct* behavior [140].

5 The inability to make a decision due to information overload or overthinking is called *analysis paralysis* [140].

One might say that decision-making is not a particularly novel topic in visualization research. This holds true for the broad area of decision-making that comprises a wide range of scenarios and disciplines. However, the task of finding the most preferred solution among a number of alternatives has not yet received specific consideration. While visualization research has presented a number of decision-making design studies and approaches that help users define "goodness", the focus is still on helping users analyze data sets rather than specifically supporting choice tasks that require trade-offs. Empirical investigation of choices, e.g., cognitive biases or evaluation metrics, are often based on controlled artificial data and simplified decision tasks that do not necessarily reflect real-world choices.

In this thesis, we investigate multi-attribute choices in the wild using well-known design study methodologies, thus moving from artificial settings towards real-world settings. While functions defining the importance or aggregation of attribute values simplify the decision process, dynamically evolving preference information calls for visualization designs that equally support making trade-offs among options and attributes. This also applies to choices that involve mutual effects among intermediate decisions. Where existing visualization design studies mainly address elementary analytic tasks like value retrieval or correlation analysis as a precondition for informed decisions, we consider the choice of a preferred alternative as a task on its own [67]. This particularly means to take into account what the human decision-maker brings to the table, such as prior knowledge, decision strategies, or irrationality. For this purpose, we borrow from disciplines like cognitive science or decision theory, which offer a long history of studying how humans think and make decisions. Finally, the effectiveness of visualizations needs to be studied on real-world choice problems and on people who are, in fact, decision-makers. Given the significance but also the difficulties of multi-attribute choices, this dissertation aims to shed light on what makes a visualization explicitly support choice tasks rather than analytic tasks.

1.2 RESEARCH CHALLENGES AND LEADING QUESTIONS

Multi-attribute choice is the task of *"finding the best alternative among a fixed set of alternatives, where alternatives are defined across several attributes"* [64]. What exactly is it that makes choices difficult?

We consider multi-attribute choice to belong to the class of *constructive* [275] or *wicked* [223] problems. Both center around the idea that the problem and the solutions to it co-evolve together, with progress towards constructing one affecting the progress of constructing the other [223, 275]. For wicked problems, Rittel and Webber identified ten characteristic properties that hint at the challenges posed by such prob-

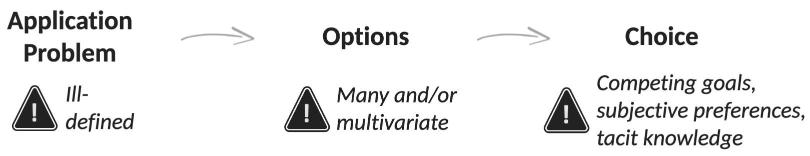

Figure 1.2: What makes choices difficult?

lems. We revisit a selection of these properties against the background of multi-attribute choice:

- **No enumerable set of potential solutions** – Solutions to an ill-defined problem can be manifold. Large sets of options with slightly varying properties are usually considered (Figure 1.2, center). Computing has made their generation fairly easy, even if options come with many properties. As any new idea about a desired outcome might become a solution candidate, it is a matter of judgment when to stop enlarging the solution set. In any case, decision-makers need to work through large amounts of data to gain insights that shift their choice in one or the other direction. This process suffers from the well-known issues associated with large data volumes and complex relationships.
- **No true-or-false solutions** – Although desired and undesired outcomes might be partially formulated as preferences, there is no formal definition of correctness. Consequently, solutions cannot be meaningfully right or wrong, only satisfying or not. From cognitive science, we know that many decisions involve affective and intuitive evaluations rather than pure rationality [25]. Whether a trade-off is considered a good choice is highly influenced by the experience, prior knowledge, and intuition of the decision-maker, which are often difficult to grasp (Figure 1.2, right). They might even shift as decision-makers learn what level of performance is achievable under which conditions. Interactive visualizations could serve as an environment for decision-makers to discover what is important to their choice and to "*leave traces that serve as self-cues*" [127].
- **No stopping rules** – As there is no single optimal solution, it is difficult to tell when a choice problem is (sufficiently) solved. Decision-makers can always invest additional effort to find a better solution. Choice problems are ill-defined in the sense that desired outcomes cannot be described by definitive conditions (Figure 1.2, left). Decision-makers might not know what they are looking for at the start of their search. Demands on a solution can be incomplete, contradictory, and evolving. When choosing a car to buy, decision-makers might be willing to increase the budget, once they realize that there are better warranties or gas mileage available than they had previously supposed. Often, solving

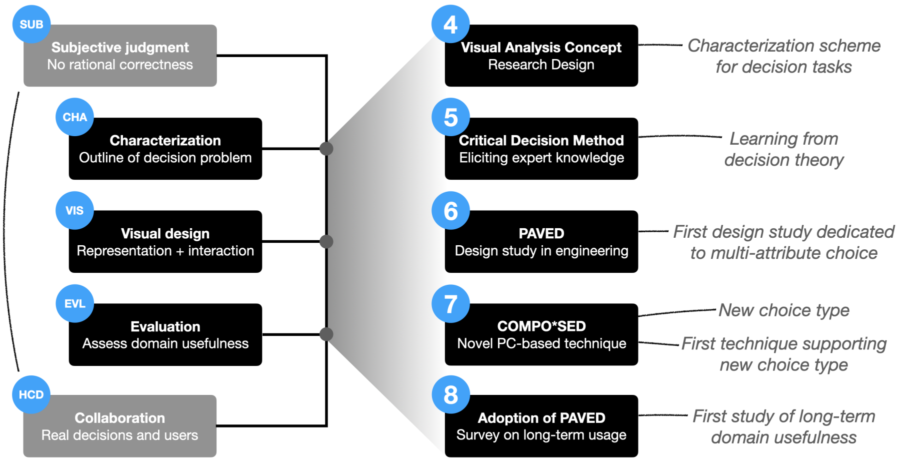

Figure 1.3: Research challenges (left) that are addressed by the contributions presented in this thesis (right). The overarching challenges C_{SUB} and C_{HCD} generally influenced the concepts and techniques.

one aspect of the problem (e.g., improving in one criterion) introduces another problem (e.g., impairing another criterion).

- **No transferability** – There can always be a property that overrides a choice's commonalities with choices already made, such that no two choices are identical. Despite their resemblance, solutions, requirements, or even the decision-maker's world view might have changed. Consider replacing an old digital camera with a new model after a few years. The range of available cameras might be different now and prices or materials of a camera might have changed. The decision-maker might now prefer a light-weight compact camera over a fully-equipped professional camera, shifting the choice into an entirely new direction. Choices can only be understood in the context of their making and dedicated visualization support needs to be carefully developed and adopted on a case-by-case basis.

What impacts for visualization research emerge from these characteristics? Visualization research aiming to support data-informed decision-making faces the following challenges (Figure 1.3, left):

C_{SUB} **Subjective judgment of satisfaction** rather than objective correctness is the basis of any decision. Subjective judgment is constructed from a decision-maker's intuition, tacit knowledge, and (domain) experience. It shapes the interpretation of the data. In turn, it typically evolves by iterating over the data as decision-makers incrementally learn about the problem, e.g., what level of performance is achievable under different conditions. Data comprehension and preference construction thus influence each other. To consider this interplay in a visualization approach, human-centered design is needed.

C_{CHA} **Characterization of decision-making activities** is essential to inform the design of effective visualizations. The understanding and description of a targeted decision problem guides all subsequent steps in visualization design and validation. It also contributes to generalizing lessons learned beyond individual cases. We do know that real-world decision-making goes beyond data comprehension. But visualization research has made few dedicated efforts to define decision problems [195]. It lacks a structured overview of what does (not) make a decision task. As long as decision tasks are not well understood, research claims in the direction of decision support are likely to remain vague.

C_{VIS} **Visual design and interaction** need to provide simple means for decision-makers to investigate relationships among attributes and options, incrementally construct preferences, filter options into acceptable and unacceptable regions, or compare the gains and losses associated with a trade-off. Visualization research has rarely provided tools that are designed to compare decision alternatives and make a final choice [195]. Simplicity and direct retrieval of decision elements might be difficult to reconcile with the need for scalability that can result from many options to navigate or many attributes to balance.

C_{EVL} **Evaluation of decision support** should aim for feedback about a tool's usefulness for decision tasks in the wild. Case studies should ask target users to make their own real-world decisions rather than work through a protocol prescribed by the visualization researchers. Methods and approaches to evaluating a visualization's ability to support decision tasks are still in their infancy [67]. Ideally, evaluations study the long-term, self-initiated use of the tool in the targeted domain. Such an endeavor, however, heavily relies on an appropriate setting within a collaboration with target users (time, willingness, confidentiality, etc.).

C_{HCD} **Collaboration with decision-makers** is needed to meet the subjective nature of any decision and to ensure that visualizations are designed for problems that real decision-makers actually face. Decision-making has mainly been studied on binary tasks or decisions of narrow complexity in artificial settings. Decision tasks taken from real-world settings typically come with higher complexity [67]. Visual designs need to be developed in constant consultation of decision-makers to meet the targeted problems and users where they are. Still, visualization research rarely pays attention to decision-makers and tasks early in the process.

The challenges outlined above are detailed as research gaps in Section 3.3, where we review related works. They touch different aspects of decision-making that have not yet been definitively answered by visualization research. This gives rise to four research questions that guide the concepts, methods, and visualization techniques presented in this thesis. The central research question is: **How can we design and validate interactive visualizations to effectively assist experts in making real-world choices among many multi-attribute alternatives?**

This involves the following leading questions:

RQ1 What decision strategies do experts apply in the wild? What is the role of cognition, domain knowledge, and experience?
RQ2 Can approaches from other disciplines be adopted in studying decision problems from a visualization perspective?
RQ3 How can options and consequences be effectively represented, explored, and balanced? How to support decision-makers in constructing their preferences as they learn about the problem?
RQ4 How can we assess a visualization's ability to effectively support experts in making their own real-world choices?

Experts, in this context, are characterized by specialist knowledge as well as implicit experience from repeated decisions in their field. They can also be held liable for their choices, which implies the need for justification. We provide details on these characteristics in Section 2.3. For these reasons, we focus on experts as our primary target group.

Through a detailed investigation of the above research questions, this dissertation aims to address the gap between multi-attribute choices being a common task among domain experts on the one hand and the lack of discourse on information visualizations explicitly serving decision-specific information needs on the other hand. All four research questions contribute to a better understanding of how information visualization can assist expert decisions in the wild.

1.3 CONTRIBUTIONS

We address the aforementioned leading questions with user-centered research at the intersection of data-informed decision-making and information visualization. For this purpose, we consult real decision-makers on their real-world data and decision tasks. We have access to this real-world setting within a long-term collaboration with domain experts from engineering design.

In this section, we summarize the results of our research at two levels of granularity. We first highlight the significance of our work by describing how we expect our results to affect the way visualization research addresses decision-making at a high level. We then describe our individual contributions in more detail.

We expect the following high-level implications to result from the findings presented in this thesis:

- From a **theoretical point of view**, we expect to broaden the current visualization theory by a characterization of decision-making activities in general and multi-attribute choice tasks in particular. We anticipate that synthesizing our experience in studying real-world decisions with decision models and theories from human science will help us carve out the subjective, context-based, and constructive nature of expert decisions. By providing a concise yet descriptive vocabulary, we expect to make the design space of decision tasks more tangible for researchers who aim to define their decision support in distinction to other works.

- From a **methodological point of view**, we expect findings from other disciplines to help us refine the design study methodology for decision problems. In particular, we intend to adopt methods from human science to establish a decision-focused way of conducting problem characterizations. More concretely, we expect to provide methodological prescriptions for eliciting expert knowledge underlying decision-making. Through applying the design study methodology, we anticipate to gain experience with effective ways of decision-focused visualization design. This also includes the collaborative aspects of user-centered design. We hope that our lessons learned from working with real decision-makers will transfer to other domains and provide guidance for future advances in problem-driven visualization research.

- From a **technical point of view**, we expect to develop visualization tools in close collaboration with domain experts that help them solve their application-specific decision problems. The design artifacts will contribute to the ongoing discourse on how decision-making is a major goal of visualization research. They will add to the body of visual encodings and interaction techniques that advance decision-related activities on large data sets, such as incrementally constructing preferences, comparing gains and losses, and reconciling conflicting information for a final choice. From observational and interview studies in real-world decision settings, also studying adoption, we expect to obtain useful evidence regarding how well visualization idioms support decision-making activities in particular contexts, and how broadly they might apply.

We show that the consideration of prior knowledge, experience, and decision strategies of a decision-maker plays an important role in assisting real-world decisions with visualizations. In summary, this dissertation advances the state of the art through the following scientific contributions (Figure 1.3, right):

- **Theoretical considerations** help visualization researchers precise their decision support claims (**$\mathbf{RQ_1}$**). We contextualize multi-attribute choice as a constructive decision problem that is situated in the scope of data science, human science, and information visualization. We also propose a characterization scheme that abstracts the diversity of decision tasks using pairs of data, user, and task properties. We instantiate this scheme to provide an operational definition of the choice task targeted in this thesis. Finally, with co-dependent choices, we propose a novel type of constructive decision problem.

- **Methodological guidance** supports eliciting the role of decision strategies, domain knowledge, and expertise in real-world choices (**RQ1**). We adopt an interview technique from cognitive science to systematically study these cognitive foundations (**RQ2**). Our main contribution is a feasibility study that indicates its applicability also in decision contexts involving large data and technology artifacts. This marks the first step towards learning from other disciplines to establish prescriptions for decision-focused problem characterizations in design studies.

- **Technical visualization artifacts** support experts in making individual as well as co-dependent multi-attribute choices (**RQ3**). PAVED is an interactive parallel coordinates visualization to ease cost-benefit trade-offs. It is the first design study dedicated to multi-attribute choice. A lossless yet compact overview of alternatives and simple interaction enable decision-makers to construct preferences as they learn what level of performance is achievable. We extended the design principle to COMPO*SED, a novel technique that, for the first time, enables decision-makers to extend their trade-off analysis to choices that affect each other. It explicitly visualizes side effects as the bottleneck of co-dependent choices. Three cascading interaction mechanisms help decision-makers subordinate individual goals in favor of those of the overall decision. The domain usefulness of both tools has been confirmed in observational studies and usage scenarios on real-world decision tasks (**RQ4**). Besides that, we surveyed the long-term usage of PAVED for decision support in the experts' work environment after four years. Four out of ten domain experts still use PAVED on their own initiative for its compact overview and reduced yet effective filtering mechanisms.

1.4 STRUCTURE OF THE THESIS

The remainder of this thesis is organized as follows.

Part II Part II provides an overview of discourses in human science, real-world applications, and information visualization.

Chapter 2 (Multi-Attribute Choice) summarizes understandings of tasks and models from decision theory, characterizes decision-makers as users, and introduces engineering design as an exemplary application for professional decision-making.

Chapter 3 (State of the Art) reviews background work in information visualization, covering methodologies for problem-driven visualization research and visualizations relevant in the context of decision-making. Finally, it outlines the derived research gaps.

Part III Part III provides the concept of this thesis and demonstrates its application to different visualization challenges raised by multi-attribute choice.

Chapter 4 (Concept) provides the research scope, goals, and targeted multi-attribute choice task. On this basis, it presents the research design of this thesis.

Chapter 5 (Knowledge Elicitation) proposes an interview technique from cognitive science to address the lack of prescriptive steps for eliciting domain knowledge in visualization research.

Chapter 6 (PAVED) presents the results of a design study in the field of engineering design to address the lack of visualization solutions for real-world multi-attribute choices.

Chapter 7 (COMPO*SED) proposes a novel visualization technique to address the lack of support for co-dependent multi-attribute choices.

Chapter 8 (Long-Term Adoption) investigates the adoption of a visualization solution in the wild after four years to address the lack of longitudinal studies of decision support.

Part IV Part IV summarizes the lessons learned of this thesis and suggests directions for future research.

Chapter 9 (Summary) distills the main findings, limitations, and contributions of this thesis.

Chapter 10 (Outlook) discusses research topics that are out of the scope of this work but are relevant to pursue in the future.

Part II

BACKGROUND

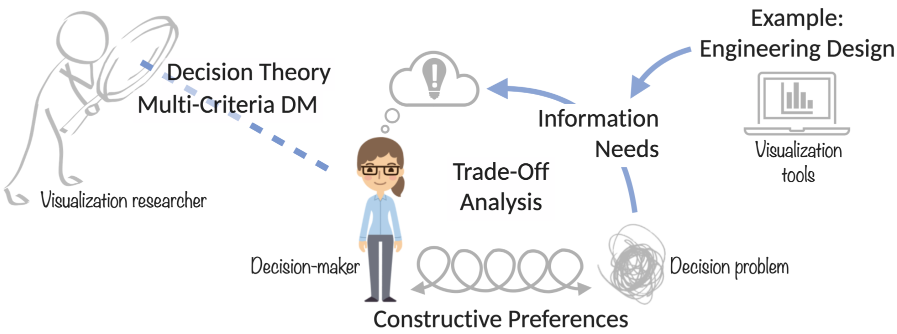

Figure 2.1: This chapter views multi-attribute choice through the lens of related disciplines to inform the research of visualization support.

2 MULTI-ATTRIBUTE CHOICE

We cannot escape making choices in our personal and professional lives. Sometimes choices are ordinary, like whether to have coffee or tea in the morning. Others are more critical and require more careful balancing of consequences. Decision problems have been studied for a long time in various disciplines that take different perspectives. We reflect on which of their concepts could help us understand and research multi-attribute choices from a visualization perspective (Figure 2.1). While the findings from these disciplines are likely useful for visualization research, at some point, we will need to leave that track to study decisions in the context of large data and technology artefacts [67].

What the reader can expect from this chapter:

- A review of decision definitions, which identifies four recurring components that hint at a cross-domain understanding (Sec. 2.1).
- A review of decision theory branches, showing that constructive models are most compatible with our targeted task (Sec. 2.1).
- A review of multi-criteria decision-making concepts, which frames multi-attribute choice as an a posteriori problem relying on a trade-off analysis among Pareto-optimal options (Sec. 2.2).
- A summary of the information needs of expert decision-makers as opposed to casual decision-makers and analysts (Sec. 2.3).
- A characterization of engineering design as a decision-centered application domain and a literature review of domain-specific yet inspiring visualizations (Sec. 2.4).

2.1 DECISION THEORY

Decision theory is an interdisciplinary topic that is concerned with the reasoning underlying people's choices. Researchers from, e.g, economics, behavioral science, management, consumer research, psychology, or cognitive science, contribute to a rich history of studying decision-making among individuals. As such, the topic involves many views and concepts. After presenting an overview of different understandings of decision tasks (Section 2.1.1), we will focus on three branches: *normative* decision theory, which attributes full (economic) rationality to the decision-maker (Section 2.1.2), *descriptive* decision theory, which assumes bounded rationality (Section 2.1.3) and thus builds upon an observation of real decision-makers (Section 2.1.4), and *constructive* decision theory, which emphasizes the continuous adjustment of heuristics during the decision process (Section 2.1.5).

2.1.1 *Decision Task*

Throughout the progression of visualization research, making a decision has been considered the ultimate goal of a visual analysis session [30, 267, 289, 290]. Existing abstractions of visualization tasks [38], cognitive biases [65], and knowledge generation models [229] connect to decision-making as a high-level process, but do not provide a characterization of the decision task itself. In general, there is no universally agreed definition for decision as a task, and the term is used in various contexts and meanings. One reason might be that decisions are studied across a variety of disciplines and establishing a unifying foundation that fits all fields is highly challenging [1]. The lack of a clear definition can make it difficult to design or evaluate visualizations aiming at decision support. In particular, it complicates ambitions to abstract the problem-driven research beyond isolated studies to develop guidance on how to design for choice tasks, which could help operationalize experience gained through design studies. Dimara and Stasko have made a first move towards leveraging guidance from decision theory by investigating to what extent decision tasks are evident in visualization research [67]. While they elaborate on what makes a decision a user task in the context of visualization, they do not expose the individual notions of decision tasks as they have been expressed in other disciplines.

How decisions or similar problems are understood in other disciplines could help understand and refine the characteristics of decision tasks for visualization. In this section, we contribute a collection of definitions from different disciplines that maps the landscape of existing high-level understandings of decision-related tasks (Table 2.1). We aim to carve out their shared understanding and distill the major ingredients. The overall goal is to gain an overview of the spectrum of decision tasks, indicate their relationship to analytic visualization tasks, and ultimately pave the way for an understanding that can serve as a foundation for our visualization design goals.

Ingredients of A Decision Task

Exploring the commonalities among the definitions, we identified four recurring components: decision activities, alternatives, attributes, and goodness indicators (Table 2.1, color-coding). The definitions reveal mandatory components, namely a decision activity (Table 2.1, blue) and the alternatives among which the decision has to be made (Table 2.1, yellow). Activity-wise, the definitions refer to *choosing* (7 out of 14 definitions), *selecting* (4 out of 14), or *finding* (3 out of 14) an entity as the primary responsibility of the decision-maker. *Choose* and *select* have an almost identical meaning and can be used interchangeably in most situations. *Find* carries a slightly different notion, namely that of a search task. According to Brehmer's and Munzner's typology of abstract visualization tasks, searching for targets with particular characteristics, like a compromise, entails browse or explore tasks [38].

While the presence of alternatives to choose from is always mentioned in the definitions, some characteristics are under-specified: the targeted size of the alternative set and whether it is fixed and known in advance. The set of options to choose from can vary in size from one definition to the next – if explicitly specified at all. 10 out of 14 definitions do not specify the magnitude of the alternative set. In the remaining definitions, the specifications range from *two or more* [59] over *more than two* [41] and *several* [220] to a *larger set* of alternatives [211]. Although these specifications mainly imply smallish numbers and none of the presented definitions names *many* alternatives as the targeted magnitude, the size of the alternative set is theoretically unbound. An explanation for the missing size specifications might be the general perception that a choice naturally involves multiple alternatives. Although Luce observes that most choice theories assume a well-defined set of alternatives [154], only three definitions explicitly state that the set of available alternatives is *fixed* [64, 196] or *known* in advance [79, 196]. Such problems fall into the category of multi-attribute decision-making, whereas in multi-objective optimization, the alternatives are not necessarily finite or known in advance [175].

Some components of a decision task are not always made explicit in the definitions: attributes that represent properties of the choice alternatives (Table 2.1, purple) and goodness indicators that help evaluate and compare the desirability of alternatives (Table 2.1, green). Only half of the definitions (6 out of 14) mention the attributes across which alternatives are defined (Table 2.1, purple). Referring to *multiple* [79, 120, 202], *several* [64], and *a set* of attributes [256], the descriptions use variations of *two or more* [41] to specify the number of attributes that characterize each alternative. Again, an explanation for this being an optional detail in a definition might be the general perception that

Table 2.1: Definitions from different fields to establish a cross-disciplinary understanding of decision tasks by comparing the mentioned decision activity, alternatives, attributes, and goodness indicators. The definitions are grouped by their tag coverage, with the richest descriptions at the end of the list.

Year	Definition	Discipline	Ref.
-	"the cognitive process of choosing between two or more alternatives"	Psychology	[59]
2015	"an individual selects one option from a larger set of alternatives"	Psychology	[211]
2001	"choose among a set of multiattribute alternatives"	Psychology	[202]
2019	"choose among more than two options that are characterized by two or more attributes"	Cognitive Science	[41]
1976	"select that alternative, among those available, which will lead to the most complete achievement of your goals"	Economics	[253]
1993	"the decision maker [...] surveys a known and fixed set of alternatives, weighs the likely consequences of choosing each, and makes a choice"	Psychology	[196]
1997	"finding the most attractive stocks [...] in a portfolio, which is essentially a matter of arbitration between the risk and the return"	Finance	[106]
1998	"choosing alternatives based on the values and preferences of the decision maker"		[90]
2006	"selecting the best option out of a set of alternatives"	Human-Computer Interaction	[24]
2019	"choice from several alternatives on the basis of a subjective value"	Cognitive Science	[220]
2001	"Given a collection of objects, each described by the values associated with a set of attributes, find the most acceptable such object"	Information Visualization	[256]
2014	"selection of the best alternative from pre-specified alternatives described in terms of multiple attributes"	Computer Science	[79]
2017	"finding the best alternative among a fixed set of alternatives, where alternatives are defined across several attributes"	Information Visualization	[64]
2023	"the choice of a subset of projects [...] with the aim of maximizing the value of the portfolio with regard to multiple [...] criteria"	Operations Research	[120]

a choice naturally involves multiple, sometimes conflicting, properties that need to be balanced.

With two or more attributes influencing the decision among alternatives, a single optimal solution does typically not exist. The majority of definitions (10 out of 14) mention in what respect the alternatives are evaluated (Table 2.1, green). However, the indication of goodness remains fairly vague for many of them, i.e., referring to the *consequences* [196] of a choice or the *best* [24, 64, 79] or *most acceptable* [256] solution to be chosen. Other definitions provide slightly more details on the underlying goodness indicators, e.g., *subjective value* [220], *most complete achievement of one's goals* [253], *values and preferences of the decision-maker* [90], or, in the case of operations and finance, the *value of the portfolio* [120] or *arbitration between the risk and the return* [106]. All of these notions align with the general understanding that, with multiple attributes to consider, goodness cannot be uniquely defined and depends on the context of the decision task.

To summarize, although different views on decision-making exist among several disciplines, we realize that these differences mainly refer to how humans should, could, or do make decisions (we will detail this in Section 2.1). In contrast, a cross-domain understanding of a decision task itself appears to be feasible. While our collection of definitions is not meant to be exhaustive and might be subject to selection bias, the comparison indicates that the range of interpretations is not as wide as one might think.

Relation To Analytic Tasks

Knowing the user task to design for is essential for visualization researchers in choosing or building visual representations that help humans solve problems [235]. Problems and user tasks can be described in the terminology of the target domain or mapped to a generic, domain-agnostic description (often based on a task taxonomy) [185]. Dimara and Stasko found that, while decision-making is emphasized as a core goal, visualization research currently centers around analytic tasks [67]. Visualizations typically serve one of three major goals: a) search for patterns and trends that feed hypotheses in an *exploratory analysis*, b) confirm or reject hypotheses in a *confirmatory analysis*, or c) communicate confirmed results for *presentation* purposes [235]. Further high-level tasks discussed in information visualization include sense-making [210], insight generation [267], or the identification of cause-effect relationships [7]. In the process of pursuing these high-level tasks, a user might engage in basic activities like retrieving a value, filtering, finding extrema or ranges, sorting, finding anomalies, clustering, identifying correlations [8], or comparing [38]. These are called low-level analytic tasks [8].

Data comprehension fostered by elementary analytic tasks is likely an important precondition for making decisions, because informed decisions require a good understanding of the alternatives and attributes involved [64]. However, there is more to making a good decision than performing well on elementary analytic tasks. Making data-based decisions differs from data analysis in that it serves different user goals and is subject to different conditions. The goal of a decision task is not to retrieve values, determine extrema, or find anomalies. Rather, it is to combine relevant information in a meaningful way so as to finally choose the most preferred option among several possible alternatives [67]. Decisions require an understanding of patterns in the data to be synthesized with subjective judgments of decision-makers. Which options are preferred or why is difficult to capture in well-defined goodness metrics (e.g., weights), which would be necessary to process them from a purely analytic perspective. Despite a proper understanding of the relevant data, people might still make irrational decisions as a result of limits in human reasoning (compare Section 2.1.3) or biases in decision tasks [65]. Consequently, Dimara and Stasko advocate for a visualization-focused formalization of relevant high-level tasks such as decision-making and a better understanding of their relationships, overlaps, and differences [67].

With their research at the intersection of multi-objective optimization and visual analytics, Hakanen et al. take a first step towards such a clarification of tasks [88]. From the general workflow of multi-objective optimization, they extracted seven high-level tasks contributing to a decision task. The authors then selected low-level analytic tasks from three taxonomies in visualization literature [94, 124, 303] that facilitate each of these high-level tasks (Figure 2.2). A follow-up work generalized the high-level tasks and further derived requirements to inform the design of a visual analytics system that has been validated in comparison with an interactive multi-objective optimization framework [89]. While these task analyses target interactive decision-making approaches, where preference articulation and generation of solutions alternate (details are given in Section 2.2.1), it can still provide inspiration for similar endeavors that address the characterization of tasks for other types of decision-making processes.

2.1.2 *Normative Decision Models*

Early decision-making models traditionally assumed a rational economic decision-maker who processed information in a computer-like way [157]. This branch of decision theory is called normative because it describes how people *should* make decisions with ideal economic behavior. Normative models postulate that decision-makers have a set of well-defined preferences and, on this basis, that they are skilled to calculate with accuracy which option maximizes their subjective value, i.e., satisfaction. In processing the available information, decision-makers are assumed to follow formal rules that, according to the

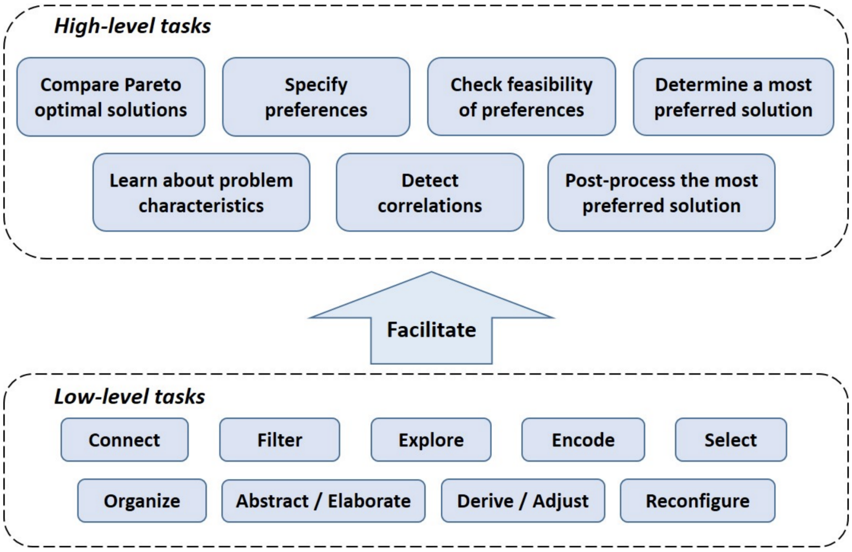

Figure 2.2: Low-level tasks facilitating high-level tasks in interactive optimization processes as introduced by Hakanen et al. [88].

theory, embody rational decision behavior. In this way, they form preference of one or the other option to arrive at a final decision.

This view manifests in one of the most well-known normative decision models: the *expected utility theory* [189]. It measures the goodness of an option with a metric, called *expected utility*, that is valid also under conditions of uncertainty. The model boils down to the decision rule of choosing the option with the highest expected utility. The expected utility of an option is made up of the utilities of its possible outcomes (i.e., how much they are preferred) and the probabilities that these outcomes occur (i.e., how certain they are). It is defined as the sum of the utilities of all possible outcomes weighted by their probability. Decision-makers strive to maximize their individually defined expected utility. Still, the theory implies that the preferences of all decision-makers satisfy the same principles of rationality. These include *completeness*, where decision-makers prefer option A over B or prefer B over A or are indifferent, *transitivity*, where preferring option A over B and B over C means to prefer A over C, *dominance*, where options that are better on at least one attribute and at least as good on all other attributes are preferred, and *invariance*, where preferences are not affected by the way options are ordered or presented. In reality, however, people have been observed to violate these principles in their actual decision-making [278].

2.1.3 *Bounded Rationality*

In the 1950s, the rationality of decision-makers as assumed by normative models was challenged by the theory of *bounded rationality* by Herbert A. Simon [251]. This theory deals with the idea that reasoning

Figure 2.3: Three-stage decision process according to Herbert Simon [252].

of humans is not necessarily consistent with rationality due to limits on their knowledge and computational capabilities. Simon argues that fully rational decisions are often not feasible in practice, because the vast amount of information needed for an exhaustive cost-benefit analysis of all possible alternatives collides with the limited time, knowledge, and cognitive resources of humans. He proposed to replace the assumption of a perfectly rational economic decision-maker with a conception of rational behavior that is compatible with the cognitive limitations of decision-makers [250]. Consequently, decision-makers are understood as humans who, in the presence of cognitive limitations, satisfice rather than maximize, i.e., they seek a satisfactory rather than the optimal decision [251].

To describe how humans actually make choices, Simon came up with a model in 1960 (Figure 2.3), which continues to influence various disciplines like computer science, economics, and psychology [252]. This model describes decision-making as a process with three stages:

- **Intelligence** – A decision problem or opportunity, i.e., an occasion that calls for a decision, is identified.
- **Design** – Possible courses of action are invented and their suitability as solutions is analyzed.
- **Choice** – The alternatives are compared and a particular course is selected. This stage is closely related to our targeted task of multi-attribute choice.

Sometimes, Simon's model is extended by a fourth stage, in which past choices are reviewed. Simon clarifies that his model is a simplified abstraction of real-world decision-making and that addressing complex decisions might involve going back and forth between the stages in view of unsatisfactory solutions or insufficient information. Still, he claims that basically any real-world decision can be characterized with the three stages.

2.1.4 *Descriptive Decision Models*

Simon's satisficing principle and three-stage model are examples for descriptive decision theory, which studies how people make decisions under natural conditions. As such, descriptive decision models are informed by empirical observations of real-world behavior. These observations have shown that, contrary to the assumption of humans possessing full rationality, in reality people might process information in unexpected ways. Tversky and Kahneman have demonstrated that people evaluate the same options differently depending on the formulation of the decision problem [278]. Their experiments showed that decision-makers behave risk-averse when possible outcomes are framed as gains (i.e., they choose the certain gain over a risky prospect of the same expected value) and risk-taking when options are framed as losses. This pattern is incompatible with the principles of rationality that underlie normative decision models, in particular with the invariance of preferences to how alternatives are presented. In response, Tversky and Kahnemann formulated their *prospect theory*, which describes how human decision-makers evaluate their gain and loss perspectives in an asymmetric manner (because the impact of a loss is felt more strongly than that of a comparable gain).

While the prospect theory centers around how people judge utility relative to a certain reference point, Payne et al. propose a framework to describe how people flexibly use a variety of strategies in response to different preferential choice tasks among multi-attribute alternatives [203]. A prior study already revealed that the same decision-maker switched to a different choice strategy as a result of varied task complexity (i.e., a varied number of alternatives or attributes) [201]. Similarly, Sunstein and Thaler suggested that the environment in which a choice is presented, e.g., the way of presenting options, attributes, and defaults, can have a significant impact on the decision made [265].

Decision Strategies

Various decision strategies have been identified and described in a number of disciplines. Each strategy basically illustrates a certain way of searching through the solution space of a decision problem. We provide a brief overview of decision strategies and their properties.

A common decision strategy is the *weighted additive rule* (WADD), where all attributes are weighted by their importance and the alternative with the highest weighted value sum is chosen [203]. This strategy is compensatory in the sense that the superior value of one attribute can compensate for an inferior value of another attribute. Compensatory strategies thus involve explicit trade-offs. WADD is one of the most demanding strategies, because it is based on a high amount of processing: the decision-maker needs to explicitly evaluate all alternatives and consider all relevant information (i.e., attribute values and importance). Due to its substantial processing of information and compensatory characteristic, this strategy is often viewed as a normative choice procedure. A simplified variant of the weighted additive rule strategy is the *equal weight* strategy (EQW), which sums the attribute values of an option without considering weights of importance [203]. Alternatively to the (weighted) value sum, the utility of an alternative can also be framed as its *frequency of good and bad features* (FRQ) [203]. The decision-maker assigns thresholds to divide attribute values into good and bad features and then counts how many of these features each alternative exhibits. Depending on the comparison of good, bad, or both features, different alternatives might be preferred.

Table 2.2: Properties of decision strategies [203] and their spread in current visualization tools [269].

	Attribute trade-off	Information ignored	Same amount of information per altern./attrib.	Processing	Score formed	Reasoning	#Ref.
	• Yes ○ No	• Yes ○ No	• Yes ○ No	• By altern. ○ By attrib.	• Yes ○ No	• Quantit. ○ Qualitat.	
WADD	•	○	•	•	•	•	3
EBA	○	•	○	○	○	○	4
SAT	○	•	○	•	○	○	0
LEX	○	•	○	○	○	○	5
MCD	•	•	•	○	•	•	2
FRQ	•	•	•	•	•	•	1

WADD = weighted additive, EBA = elimination by aspects, SAT = satisficing, LEX = lexicographic, MCD = majority of confirming dimensions, FRQ = frequency of good and bad features

Like EQW involving sums and WADD involving multiplications, the count underlying the FRQ strategy represents quantitative reasoning. Rather than a simultaneous evaluation of all alternatives, the *majority of confirming dimensions* strategy (MCD) compares alternatives in a pairwise manner [203]. The values of two alternatives are compared attribute for attribute and the alternative with a majority of better attribute values is retained. It is then compared to the next alternative, until all alternatives have been evaluated and a final choice remains.

Compensatory strategies confront conflicts among decision criteria, whereas non-compensatory strategies avoid them. One of the oldest non-compensatory strategy is *satisficing* (SAT) [250]: Alternatives are considered by order of appearance and the first alternative to satisfy predefined thresholds for all of its attribute values is chosen. If no alternative passes all thresholds, the process can be repeated with relaxed thresholds or an alternative is chosen randomly [203]. The *lexicographic* strategy (LEX) is also non-compensatory in that a decision-maker identifies the most important attribute and chooses the alternative with the best value on this attribute [203]. If two or more alternatives are on par, the procedure is repeated with the next most important attribute. The core elements of SAT and LEX are combined in the *elimination by aspects* strategy (EBA) [277]: Attributes are considered by order of importance (like in LEX) and all alternatives that do not meet the threshold on the current attribute are eliminated (like in SAT) until one alternative is left [203]. As this involves only simple comparisons of values, this strategy is based on qualitative reasoning.

Table 2.2 summarizes how the strategies compare to each other with respect to their properties. Whether alternatives are evaluated across or within attributes (4th column) differs among the compensatory and non-compensatory strategies. Attribute-based strategies are assumed to be cognitively easier to perform [228]. Non-compensatory strategies do not require decision-makers to engage in explicit trade-offs (1st column), which might help avoid emotional discomfort [102]. They can also be said to be less demanding, because they use only part of the information available (2nd and 3rd column), do not require the formation of a utility score for each alternative (5th column), and involve simple value comparisons rather than mathematical operations (6th column). On the other hand, they might overlook information that is relevant to the decision-maker. One solution is to use them for initial elimination of poor alternatives in combination with a more detailed evaluation of the remaining small number of alternatives.

The decision strategies described above have also inspired visualization research. In an attempt to understand how these strategies are supported by visualization tools, Torsney-Weir and colleagues classified 21 design study papers according to the main strategy described in the task analysis, user characterization, or case study [269]. 15 of the papers involved some sort of strategy, with the compensatory frequency of good and bad features strategy and the non-compensatory lexicographic strategy being the most common ones, followed by elimination by aspects, weighted additive rule, and majority of confirming dimensions (Table 2.2, right). Their review suggests that the satisficing strategy does not play a role in visualization research. Most of those strategies supported by visualization tools involved a consideration of one attribute at a time and a filtering of alternatives.

2.1.5 *Constructive Decision Models*

Normative (Section 2.1.2) and descriptive (Section 2.1.4) decision models aim to describe how people should or do make decisions in general, independently from the individual decision-maker. How individual decision-makers approach a choice problem can be distinguished by the way they use choice heuristics. There are basically two branches emphasizing the decision-maker's subjectivity: the *prescriptive* mechanism, which assumes a set of rules stored in memory that can be retrieved and applied when needed, and the *constructive* mechanism, which assumes that heuristics are developed at the actual time of choice [61]. The notions of bounded rationality and limited processing capacity addressed by descriptive decision models align with the idea that preferences for alternatives are not merely applied but in fact constructed during decision-making [31].

Constructive decision models are based on the idea that, in reality, decision problems are not well-defined right from the beginning but only unfold on the spot as available information are structured or restructured. In particular, this means that decision-makers, instead of having a repertoire of complete and well-defined preferences, develop them on the fly when they need to make a choice. Preferences are constructed from fragments that can represent beliefs about alternatives, previous evaluations, rules of thumb, rules for assigning weights (e.g., "if performances are comparable, price weighs heavier") [32]. These fragments depend on what is available in the particular choice situation and how easy various pieces of information are to process [32]. The notion of constructive preferences thus rejects the assumption that decision-makers simply recall preferences from a master list in memory [31]. This also links back to the idea of people using different decision strategies in response to varying task complexity [203] or, in this case, in response to gaining new insights about the problem structure during the course of making a decision [31]. An extensive learning process is thus inherent to constructive decision theory [275]. In other words, constructive decision problems are characterized by the need to "*construct at the same time the problem and its solution*" [275]. Similarly, wicked problems have been characterized as problems where "*the process of formulating the problem and of conceiving a solution [...] are identical*" [223].

Constructive mechanisms tend to be used when decision-makers are unexperienced with a particular choice, or when a choice is particularly difficult [32]. One reason for which people might construct preferences is that they lack the cognitive resources to form well-defined preferences for any eventuality. Another reason might be external factors that, e.g., lead people to prefer chocolate ice cream on one day but choose a different flavor on another day. Constructive preferences might also include meta goals related to a decision, such as minimizing the required cognitive effort, minimizing the experience of negative emotion (compare the discomfort associated with the need to make explicit trade-offs [102]), or maximizing the ease of justifying the decision [31]. As an implication of the constructive nature of preferences, choices among options are generally considered to be dependent on 1) the (meta) goals of the decision-maker, 2) the difficulty of the decision task, 3) the context, i.e., the utility of an alternative depends not only on its own properties but also on the properties of the remaining alternatives, 4) the way of eliciting the decision-maker's preferences, and 5) the presentation of alternatives (compare the asymmetric perception of gains and losses [278]) [31].

2.1.6 *Important Terms*

Multi-attribute choice overlaps with diverse disciplines like multi-criteria decision-making, cognitive sciences, modeling, or optimization. Each discipline prefers its own terms. To prevent ambiguities, we briefly define the terms as they are used throughout this thesis.

- A **decision problem** is the subject under investigation, about which the choice is made, e.g., the camera in the scenario of a photographer intending to buy a new camera.
- **Alternatives** are specific solutions to a decision problem, among which the choice is made. For a photographer, the alternatives might include different camera models, e.g., a Canon EOS 50 or a Nikon Z50. Alternatives can also be called options, candidates, or, in the field of genetic optimization, individuals.
- **Attributes** describe the properties that an alternative might exhibit, e.g., release year, focal length, mount type, price, weight, or resolution. In making a choice, one typically aims at alternatives with preferred properties.
- **Design parameters** are those attributes that define how each alternative is implemented in practice. For example, the appearance of a camera can be described by release year, focal length, or mount type. In principle, the values for these attributes can be freely chosen. In the domain of modeling, these attributes are typically called input or independent dimensions. In the domain of optimization, they are called decision variables and the space that they span is called the decision space.
- **Criteria** are those attributes that are considered in evaluating and comparing the goodness of alternatives. A photographer might consider the price, weight, or resolution of a camera model. All criteria values of an alternative make up its performance. In the domain of modeling, these attributes are typically called output, dependent, or target dimensions. In the domain of genetic optimization, the criteria are typically called objectives.
- **Preferences** express the decision-maker's priorities regarding attributes or alternatives. They can manifest in filters or weights that are assigned accordingly. Preferences are highly subjective and context-dependent. The preferences of the photographer might depend on her degree of professionalism, budget constraints, or usage scenarios (e.g., outdoor settings).

Problem-driven research also requires a certain flexibility regarding the use of terms. To bridge the knowledge gap [296], it might make sense to adapt to the terminology used in the targeted domain when talking to experts. For example, the engineering experts providing the application problems targeted in Chapters 6 and 7 adopted the terms "individuals" and "objectives" from genetic optimization theory. When presenting our research to the visualization community, however, we decided to use the terms "alternatives" and "attributes" for two reasons. First, these terms are more neutral, e.g., they are used in literature across domains (compare Table 2.1). Second, which attributes actually constitute criteria in a decision task depends on the subjective judgment of the decision-maker and might be unknown to the visualization researcher at the time of visualization design.

2.1.7 *Summary*

Early normative decision models attributed ideal rational-economic behavior to human decision-makers. Studies later showed that human rationality is bounded [251] and people in actual choice situations systematically violate the normative principles [278]. Descriptive models have been introduced to describe how people make decisions under natural conditions. They view human choice as an "*adaptive and constructive process rather than one of expressing existing well-defined preferences*" [266]. The shift from normative over descriptive to constructive decision models has a number of implications for (visual) decision support tools (compare Section 1.2):

- Decision-makers might need to learn what level of performance is achievable under different conditions before being able to construct preferences [266]. They tend to dive deep into an exploration of relevant information. A tool should make contextual knowledge accessible [217] and involve representations that have been carefully chosen with relevant analysis tasks in mind.
- As preferences are not readily available from the beginning, decision-makers need a flexible means to state or refine their preferences at any point during the decision process. With constructive preferences, the circumstances will have a greater impact on the resulting choice [32]. A tool should support incremental specification of preferences [217].
- The construction of preferences might be guided by relevant information raising the attention of the decision-maker at any point in the decision process. A tool should not require decision-makers to express their preferences in a rigid order [217].
- Compromises are central to the construction of preferences. A tool should hint at criteria conflicts and partial fulfillments to help decision-makers engage in trade-off analyses [217].
- Decision-makers might not be able to fully anticipate the consequences of their preferences. A tool should provide immediate feedback about the effects of their expressed preferences at any time point [266].

2.2 MULTI-CRITERIA DECISION-MAKING

While Section 2.1 on decision theory focused on human aspects of decision-making, this section looks at decision-making from a data perspective. At their core, most real-world decisions involve multiple conflicting properties to be considered simultaneously. In this sense, multi-attribute choice is closely related to multi-criteria decision-making (MCDM). The discipline of MCDM, however, covers a broader spectrum of decision tasks (e.g., infinite or varying alternative sets, ordering rather than selecting alternatives) and often involves imposed decision procedures [64]. Still, some of its underlying concepts also apply to multi-attribute choice. In the following, we highlight relevant concepts and discuss how they relate to choice tasks.

2.2.1 *Multi-Criteria Decision-Making Approaches*

As Horn points out, multi-criteria decision problems can be separated into two distinct tasks: the *search* for solutions and the multi-criteria *decision* [103]. According to Ishizaka and Nemery, different goals can be pursued by addressing the decision task [114]: 1) a single most preferred option or a small subset of incomparably good options, 2) an assignment of options to predefined groups or a ranking from best to worst, or 3) a description of the characteristics for each option. We do not focus on the latter tasks of grouping, ranking, or characterizing options, because they can be considered secondary tasks involved in making a choice. Generally, approaches to a decision problem can be classified according to the role of the decision-maker [103].

In *a priori* approaches, decision-makers articulate some preferences for the different attributes prior to the search for potential solutions (decide $\rightarrow$ search). Based on these preferences, the attributes are aggregated into a fitness function (most commonly a weighted sum). Well-known examples for visualization systems addressing this category are *LineUp* [85] and *WeightLifter* [198]. While fitness functions can simplify the decision process, establishing appropriate attribute preferences up front is often challenging. They may be vague at the beginning or evolve over time, especially if decisions are made on behalf of other stakeholders.

Thus, in *a posteriori* approaches, decision-makers articulate their preferences after the search (search $\rightarrow$ decide). Without an objective function, search algorithms can only compute a set of mathematically equally good trade-offs (known as the Pareto front, see Section 2.2.2). Determining the most preferred option among this reduced set of solutions is then left to the decision-maker's judgment. While the search is independent of the decision-maker's preferences, the approach heavily relies on the search algorithm identifying an appropriate distribution of solutions along the Pareto front.

Thus, in *interactive* approaches, decision-makers articulate their preferences progressively; search and decision tasks are performed in alternation (search $\leftrightarrow$ decide). A preliminary search conveys an initial idea of what trade-offs are possible. Decision-makers then express preferences to reduce the search space, such that the following search is restricted to this particular region of interest. An example for decision-makers guiding a search algorithm towards preferred regions with the help of a heatmap visualization is provided by Hettenhausen et al. [98]. While decision-makers can interactively steer the process by

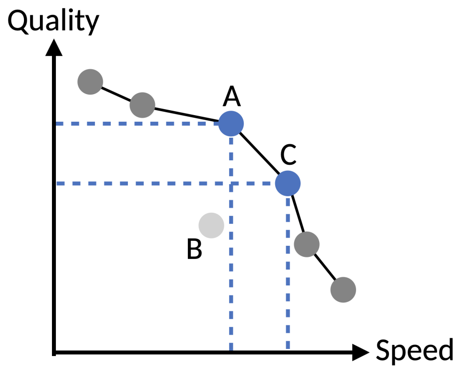

Figure 2.4: Example of a Pareto front regarding two criteria to be maximized. Point B is dominated by points A and C. Points A and C are not dominated by any other point and thus belong to the Pareto front.

providing preference information, interactive approaches have not yet found their way into real-world applications [88].

Multi-attribute choice falls into the category of a posteriori decision problems. To be more precise, the search task can be considered already solved. The resulting set of solutions, where no criterion can be improved without sacrificing at least one other criterion, forms a fixed set of alternatives among which the best alternative is to be found (compare Section 2.1.1). Multi-attribute choice can therefore be considered the decision task in an a posteriori decision problem.

2.2.2 *Pareto Optimality*

Section 2.2.1 concluded that multi-attribute choice can be considered an a posteriori decision problem, where the search has already taken place. As options are defined across multiple attributes, a unique option that simultaneously optimizes each attribute does generally not exist. When attributes cannot be optimized simultaneously, they are said to be *in conflict*. A well-known conflict in economics is that between time, cost, and quality as modeled in the project management triangle [15]. Where multiple attributes need to be considered in a decision, comparing the quality of options becomes challenging (while options can be easily ranked by how they score on individual attributes). Without any tendency regarding which attributes might be more important, such that solutions scoring better on these attributes gain superiority, we cannot generally tell whether one solution is better than another solution. However, the importance of attributes cannot be deduced from the data themselves but requires the subjective judgment of the human decision-maker.

There is only one multi-attribute case, in which we can rate the superiority of an option without any prior knowledge or judgment. If an option A does not score worse than another option B on any attribute and there is at least one attribute for which A scores strictly better than B, we know that option A is generally superior to option B. This concept is called *Pareto dominance* [58], i.e., option A dominates option B (Figure 2.4).

As we are usually dealing with large sets of options to choose from, it does not help much to know that one option dominates another option. Instead, it makes sense to restrict our choice to all those options that are not dominated at all (Figure 2.4). Let's define a multi-attribute option as a tuple $\vec{y} = (y_1, ..., y_n)$ of values $y_i \in Y_i$ for attributes $Y = \{Y_1, ..., Y_n\}$. An option $\vec{y}^*$ is called *Pareto optimal*, or non-dominated, if there does not exist another $\vec{y}$ that dominates it. In other words, no attribute of $\vec{y}^*$ can be improved in value without impairing at least one other attribute value. The entirety of all Pareto optimal options is called the *Pareto front* [58]. Without prior knowledge, no option in the Pareto front can be considered superior to another one. Thus, the Pareto front contains all options that can be considered equally good. From here, superiority of options solely depends on subjective preference information. Note that, as we are taking an a posteriori approach, the following multi-attribute choice can only be as good as the Pareto front computation (compare Section 2.2.1).

2.2.3 *Trade-Off Analysis*

Section 2.2.2 concluded that subjective judgment is needed to make a final choice among a set of non-dominated options. Thus, a *trade-off analysis* is at the core of making a multi-attribute choice. A trade-off analysis is conducted by decision-makers, who systematically explore the advantages and disadvantages of each (non-dominated) option. Another notion of trade-off analysis is the ambition to understand what trade-offs decision-makers value in certain situations, e.g., in consumer research where analysts observe consumers' choices among products whose characteristics are varied in systematic ways [118]. While the elicitation of attribute weights and trade-offs is also relevant and has been addressed in visualization research [198], we focus on the former understanding where decision-makers take the active part. Of the different strategies to search through the solution space of a decision problem (Section 2.1.4), only the compensatory strategies allow for explicit trade-offs. Trade-off analysis takes a closer look at the gains and losses associated with deciding between individual options.

When moving from one non-dominated option to another non-dominated option, attributes can be independent, harmonious, or conflicting [219]. Independent attributes do not exhibit a particular relationship (Figure 2.5a) and are equally considered in an analysis. When two attributes are in harmony, the performance of one attribute increases, too, as the other attribute is improved (Figure 2.5b). Due to this redundancy, one of them can be neglected in the following considerations. With conflicting attributes, a gain in one attribute comes with

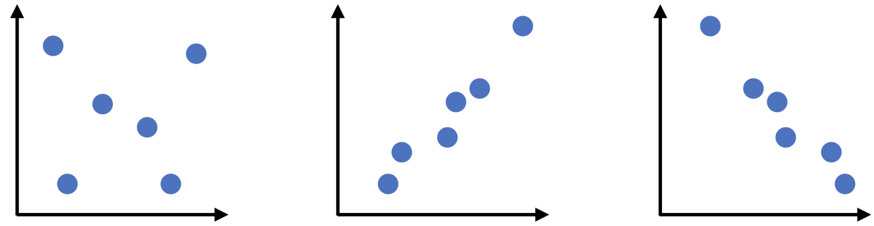

(a) Independent (b) Harmonious (c) Conflicting

Figure 2.5: When comparing the advantages and disadvantages of decision options, attributes can be independent, harmonious, or conflicting. In these examples, larger values are preferred to smaller ones.

a sacrifice in at least one other attribute [58] (Figure 2.5c). Sacrificing one attribute in return for a gain in another attribute is called a *trade-off*. Different options produce trade-offs among different attributes. In particular, an option that optimizes one attribute requires compromises in other attributes. Multi-attribute choice, however, requires decision-makers to decide for only one trade-off.

Which of the given options should they choose? To make a choice, decision-makers need to think about which sacrifice in one attribute they are willing to accept for a certain gain in another attribute: they apply subjective preferences to arrive at a compromise solution. It means to compare and balance the costs and benefits of each option by exploring the impact of certain improvements on the remaining attributes. This fosters an understanding of the interdependencies between the attributes and the feasibility of one's preferences. As a consequence, preferences might also evolve dynamically as decision-makers learn which option quality is achievable under which conditions. All of this aims at making the conflicting nature of the involved attributes more transparent. As a result, the choice is made with full comprehension of the advantages and disadvantages of each option in the context of a concrete application, i.e., it is an *informed* choice.

Depending on the number of options and attributes available, a full trade-off analysis will not be feasible, mainly due to the limited cognitive capability of our brains (Section 2.1.3). With the increasing size and complexity of solution spaces, many decision-makers turn towards visualization tools to help them understand the data on which their decisions are based. Tušar et al. found that requirements of trade-off analysis typically go beyond what traditional multivariate visualizations have to offer [276]. Based on a star plot, Trinkaus and Hanne enable trade-off analysis by allowing decision-makers to drag criterion vertices towards smaller or larger values, making the visualization update to the next closest matching option [274]. A visualized decision horizon, logging of navigation paths, and storing favorite options additionally facilitate trade-off analysis. Similar to the decision horizon, Berger et al. predict combined criteria changes in a local neighborhood when increasing or decreasing design parameter values

[28]. Still, in a study with organizational decision-makers, Dimara and colleagues identified the need for dedicated visualization support for trade-off analysis: despite the availability of visualizations, the decision-makers reported difficulties in dealing with trade-offs [68]. To help decision-makers keep track of trade-offs, Chakhchoukh et al. suggest to extend visualizations with provenance functionality [47]. In an observational study with domain experts, they investigated trade-off tasks, provenance purposes, the usage of provenance functionality, and future opportunities for in-visualization provenance.

2.3 EXPERTS, NOVICES, AND ANALYSTS

Section 2.1.1 revealed that decision tasks cannot be put on a level with analytic tasks. Besides an understanding of what particular task is to be achieved, assessing the appropriateness of a visualization also requires a consideration of the user [176]. People working with data typically fall into one of three categories: professional decision-makers, casual decision-makers, or analysts. Similar to the tasks, users who are decision-makers cannot be treated the same as users who are analysts. In the following, we review the discourse on what it means to target professional decision-makers as opposed to the needs and objectives of casual decision-makers and analysts.

We refer to *professional decision-makers* as experts, whose job it is to repeatedly make decisions within their field and who are trained to do so. *Casual decision-makers* are novices, who make data-informed decisions at some points in their personal lives. We adopt the understanding of *analysts* as people whose primary work task it is to answer questions with data [68]. Both data analysts and decision-makers need to gain an understanding of the data to perform their tasks. However, information understanding can take various forms depending on the overall aim [195]. Decision-makers differ from data analysts in that they use data to make decisions, not answer questions [67].

Visualization research has a long tradition of aiming to support professional analysts as the target user [68]. Works on casual [214] or personal [104] visualization seek to make the power of visualization accessible also to novices, who aim to obtain insights from data for personal interests or needs. As part of their research agenda, Dimara and Stasko suggest to broaden the profiles of target users from data analysts to decision-makers in both professional and casual contexts [67]. Due to their distinct goals, however, visualizations that work well for data analysts are not necessarily effective for decision-makers. To tailor a visualization design to the unique information processing requirements of decision-makers, we need to know under which conditions they make a choice. For example, Alves et al. found the personality trait of conscientiousness, i.e., the tendency to follow goals and prioritize tasks, to affect visualization-based decision-making

[6]. Similarly, Conati et al. found different user characteristics like cognitive abilities as well as visualization or domain expertise to affect performance regarding decision tasks [57]. Further conditions include the decision-makers' physical or mental environment, their preferences and goals, their knowledge and experience, their decision strategy as well as external specifications, constraints, or stakeholders.

How Are Decision-Makers Similar to Analysts?

The lowest common denominator of decision-makers and analysts is that both of them have an information need that requires them to work with data. As a result, they will start exploring the data, aiming to discover interesting trends and patterns. A decision-maker and an analyst working with data on the same topic might consider similar aspects of the problem [195]. Both might also engage in synthesizing their findings and drawing connections between patterns or insights [195]. For organizational contexts, Dimara and colleagues have found a significant overlap between both roles: decision-makers also engaged in data analyses and analysts were also involved in decision-making beyond micro-decisions in their own workflow [68]. The authors conclude from this finding that aiding decision-makers might also (partially) support analysts as a side effect.

Both decision-makers and analyst might need to report their data exploration (procedure and results) at some point [68]. This includes identifying the best way to present the data to downstream consumers (e.g., other stakeholders in the case of decision-makers or decision-makers in the case of analysts) and relating the data or findings to what one cares about [5]. Generally, both user groups value transparency in the sense of being able to explain the steps they followed to arrive at a conclusion [68]. Decision-makers rely on it to communicate and justify their decisions, while analysts rely on it to provide actionable results for their downstream users, e.g., decision-makers.

How Are Decision-Makers Different From Analysts?

Although decision-makers and analysts share some commonalities in working with data, their analysis contexts and underlying goals can largely differ. First, the different (educational) backgrounds of both user groups might affect their information processing workflows [68]. Decision-makers are specialists in their field[1], which allows them to identify domain-specific patterns in the data, focus on those that are important, and interpret them in the light of their domain-specific semantics and goals. Analysts typically have a statistical or informatics background, which qualifies them to efficiently digest whatever information seems relevant from structured or unstructured data but might not go beyond such general, domain-agnostic findings.

1 Casual decision-makers can also be considered specialists in their field, with their field being their personal environment.

Given their different backgrounds, decision-makers and analysts take different roles in a data exploration workflow. Although one group can perform its tasks without the other, there might be a transfer of results from the group of analysts to that of decision-makers. The findings that analysts generate during their data exploration might or might not be taken up by decision-makers [68]. In this sense, analysts could be seen as recommending findings that might be actionable to potential decision-makers. These roles are tightly coupled with their overall goals. Analyses of decision-makers, at least in organizations, have been found to be driven by a well-defined question (i.e., neither a hypothesis nor an exploration task) [67]. In some sense, this conveys the notion of an analysis following a directed search for some insight that might make a difference for the question. In contrast, analysts mainly aim at an undirected search, detecting unexpected relationships in the data, which might or might not be relevant to the question of a decision-maker. To summarize: decision-makers work with data to collect evidence for their following decision, while analysts do it to highlight interesting patterns [67].

Given their backgrounds and overall goals, decision-makers and analysts have different requirements for information processing. An interview study on organizational decision-making has found information overload to be an issue for decision-makers, while analysts engaged in detailed data analyses and produced verbose reports [68]. The same study found differences in the interests regarding the type of represented information, where analysts mainly dealt with the quantitative aspects of a topic, while decision-makers also relied on qualitative (sometimes informal) information, which can be difficult to integrate. Different requirements might also result from the gap between the (perceived) responsibilities of decision-makers and analysts, as reported by a participant of an online survey: "*A decision maker has to live with their decision where an analyst can just say what the best thing is and walk away!*" [68]. This could be why decision-makers might have a stronger need for uncertainty awareness, which is sometimes interpreted by analysts as if decision-makers ignore analysis results unless they depict oversimplified information [68]. The different approaches to data analyses generally hold the potential to provoke misunderstandings when decision-makers and analysts work together.

Oral et al. compare decision-makers and analysts with respect to different aspects of information processing [195]. Decision-makers often need to direct their attention and focus towards information that shift their decision in one or the other direction. In contrast, analysts may engage in an exploratory search, diving into various aspects of the data without being constrained by the need to prioritize so as to arrive at a decision. Decision-makers benefit from simplification of the information they base their decisions on, e.g., by eliminating less relevant information or breaking down complex information into categories and meaningful patterns. In contrast, analysts aim to understand the data in their full complexity to reveal hidden relationships. A core task of decision-makers is to determine the utility of a solution by weighing its risks and benefits in comparison to alternative solutions. Instead of such trade-offs, the focus of analysts is more on the data quality and the overall relevance of their findings to the broader topic. Finally, while decision-makers balance conflicting demands to arrive at a compromise solution, analysts are more interested in synthesizing findings to contribute to a broader body of knowledge.

How Are Expert Decision-Makers Different From Casual Ones?

We previously outlined differences between decision-makers and analysts regarding their approaches to information processing. To further carve out the characteristics of professional decision-makers, who repeatedly make decisions within their field, we compare them to casual decision-makers, who make decisions related to their own personal lives (e.g., the choice of a university, a place to live, a car to buy, etc.). Most problem-oriented works in visualization research target professional data analysts. Some design studies report on collaborations with expert decision-makers [35, 85, 254], but do not prioritize decision mechanisms over analytic tasks. Visualization tools specifically meant to support multi-attribute choices made by non-experts are detailed in Section 3.2.2. While they provide simple visual encodings, dynamic filtering, and mundane data sets, they do not explicitly state what distinguishes casual decision-makers from experts. As opposed to problem-driven visualization research, surveys on casual [214] or personal [104] visualization that aim to enable novices analyze data relevant to their personal lives. In the field of decision theory, expert decision-makers are contrasted with the public (referring to society) [119], which does not completely align with the idea of comparing expert and casual decision-makers. Unlike decision-makers versus analysts [195], no attempt has been made yet to explicitly contrast expert decision-makers with casual decision-makers.

The common baseline of both expert and casual decision-makers is that they process data to obtain information that shift their decision in one or the other direction. To choose alternatives over others, experts and non-experts need a means to express and apply their preferences towards (un-)desired options or properties. In professional choices, the preferences are mostly shaped by specialist knowledge and experience, while casual decisions might also involve personal taste. Generally, expert decisions have been described as large decisions in the sense of "*important [...], collaborative, or high-stakes decisions*", whereas personal decisions have been described as micro decisions [67].

Chester Barnard explicitly differentiated between decisions an individual makes as a member of an organization (we consider these as professional decisions here) and decisions the same individual makes in his personal life [19]. Herbert Simon continued this thought by contrasting decisions related to "*personal needs and results*" with decisions made "*in an impersonal sense as part of the organizational intent*" [253]. Pousman et al. note differences between visualizations targeting expert users and those targeting casual users regarding usage, data type, and insight [214]. As opposed to professional decision-makers, casual decision-makers do not have a specialist background and are not necessarily familiar with analytic thinking or reading visualizations. With casual users, the decision problems extend to life areas other than one's profession. The data underlying decisions is no longer work-motivated only but carry personal meaning such that casual users might bind more tightly to them. The type of insight generated during decision-making might also be different. While professional decision-making mostly generates analytic insights, personal decision-making might also foster reflective insights about oneself, e.g., one's beliefs and personal way of thinking.

Casual decision-makers mostly explore data about themselves [104]. As a consequence, they might be more attached to their choices, because the choices are likely associated with a long-term effect on the decision-makers life, e.g., a university that is attended over years. Emotions may generally play a greater role for casual decision-makers, potentially leading to phenomena like trade-off avoidance [153]. Professional decision-makers might be less sensitive to this bias, because they are not as closely involved. Professional choices are typically made on behalf of other stakeholders. These might be customers who do not have the required expertise themselves, or it might also be the decision-maker's employer, in whose intent a decision is made (see Simon's comparison above). Professional decision-makers might not be affected by the decision themselves (e.g., pay for the chosen item). As it is not the expert decision-makers but their stakeholders, who are affected by the choice's consequences, values like trust, transparency, and involvement are crucial.

Summary

To answer the information need of a decision-maker with visualizations, we need to understand how professional and casual decision-makers approach information processing (differently). This involves factors like educational background, evaluation criteria, (emotional) attachment, responsibility, or a guiding question. These findings provide a starting point for visualization researchers aiming to characterize the target decision-maker to design for.

Designing for expert decision-makers means to help them collect domain-specific evidence for an impersonal decision that might need to be justified in front of affected stakeholders. Expert decision-makers aim to interpret patterns in the data with respect to their domain-specific semantics and goals, which might also include qualitative information. Applying domain knowledge requires sufficient opportunities to interact with the visual representations. Decision-makers are often driven by a well-defined question. Rather than digging into all available details of the data, they prioritize those information that make a difference for the decision by eliminating less relevant information. This requires a simplification of the presented information as well as effective filtering capabilities. To arrive at a final compromise, they need to balance the (conflicting) costs and benefits of options, which places demands on their visual comparison. Expert decision-makers take an impersonal perspective on such trade-offs and act in their organization's or customer's intent. They also need to take responsibility for their actions, which often affect others. This places high demands on uncertainty awareness and transparency of the steps taken, which are essential for communicating a decision.

2.4 MULTI-ATTRIBUTE CHOICE IN ENGINEERING DESIGN

As an example for professional decision-making, this section introduces the field of engineering design. The abstraction from domain-specific details of (systems) engineering design to (co-dependent) multi-attribute choices is provided in Chapters 6 and 7.

2.4.1 *Engineering Design*

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 diverse areas of our society. Engineering design is the *"process of devising a system, component, or process to meet desired needs and specifications within constraints"* [74]. At their core, most real-world decisions pose a conflict between several criteria. Engineering design is no exception. 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. Given a performance specification, we want to know what set of input parameter values is needed. Unfortunately, there is no direct way of computing these parameter values from a desired performance specification. The absence of a direct mapping from performance to input parameters is the inherent challenge of engineering design [255, 279].

Computing has made it easy to generate large amounts of solution candidates by simulating the subject under investigation with varying design parameters. 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, engineers search for a solution that is optimal for a given application. Where single components are targeted, decision-makers thus face a *multi-attribute choice* problem.

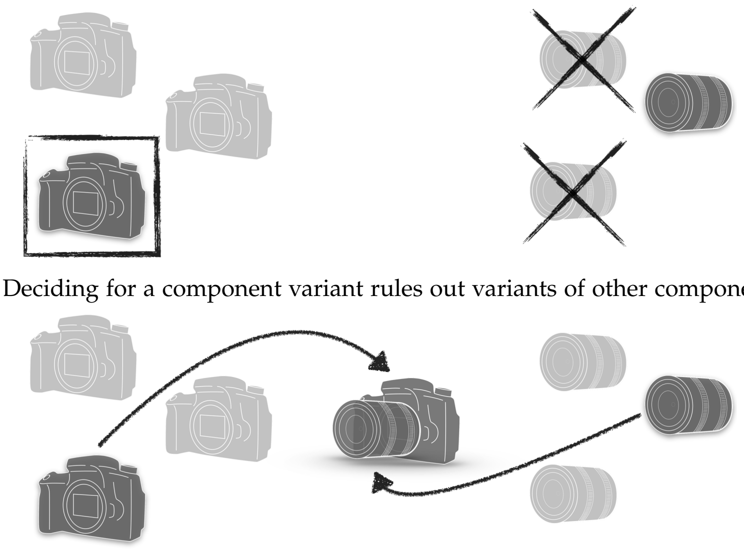

(a) Deciding for a component variant rules out variants of other components.

(b) Individual goals are subordinated in favor of those of the overall system.

Figure 2.6: Systems engineering design means to choose multiple components such that their interplay achieves a desired performance.

2.4.2 *Systems Engineering Design*

Instead of single components, many subjects to decide upon are actually systems. Many notions of what a *system* is have been expressed in the literature [137]. Most of them describe a system as a collection of components that jointly perform a function to achieve a common objective. Each component is typically designed by dedicated specialists who focus on optimizing its individual characteristics [137]. A central task in systems design is, therefore, to establish a balance or even a symbiosis [307] among the various components. With the term *systems design*, we denote the process of determining a combination of interacting components that optimizes the emergent system performance with respect to a number of objectives. Carlson-Skalak et al. introduced the term *catalog design* [44] for a two-stage process that consists of 1) specifying a system configuration [179], i.e., an arrangement of generic components, and 2) instantiating this configuration by selecting particular component variants from catalogs. We refer to the second stage, where the generic components are instantiated in a way that optimizes the system performance while the configuration does not change.

As an example, let us consider the design of a magnetic-geared motor, where motor and gear are arranged side-by-side within a common housing [161]. These motors are suited for industrial applications that require high power densities, e.g., wind energy or ship propulsion. Engineers need to decide on a motor and a gear variant such that their interplay achieves the desired outcome. For this decision, they need to be aware of two aspects: 1) deciding for a motor restricts their gear choice due to different mounting point geometries, and 2) the superiority of a motor variant depends on its interplay with the gear variants, i.e., can only be judged based on their joint performance. Similarly, deciding for a camera body restricts the lens choice due to different mount types (Figure 2.6a) and the quality of a camera body only shows in its joint performance with a lens (Figure 2.6b).

When targeted at a system rather than a single unit, deciding on suitable trade-offs among conflicting criteria thus becomes even more challenging. With a single unit, *best* solely depends on an alternative's Pareto dominance and the decision-maker's subjective preferences. When dealing with a system, the superiority of a component variant additionally depends on its interplay with the rest of the system. Making trade-offs thus extends beyond one single unit.

When facing such a task, deciding on each component independently is not an option because the chosen parts might not be interoperable. Even if they were, component-wise optimality would not guarantee a globally optimized system performance due to emergent effects. Both problems could be solved by using a multi-component model to represent the entire system. However, this means increased model complexity and computational efforts because unchanged partial simulation results cannot be reused [44]. To strike a balance between these two ends of the spectrum, an approach taking into account both 1) the variants of individual components (component level) as well as 2) their linking according to interoperability and joint performance (system level) is needed.

One approach considering both levels is an iterative optimization, where the decision-maker observes one system component at a time. However, iterative multi-attribute exploration of a system can become a tedious and, at times, frustrating trial-and-error process. In each iteration, the most preferred variant for the component under investigation is chosen (component level). Its properties are then considered in the subsequent iterations to evaluate the interoperability and overall performance (system level) when deciding on the remaining components. With an iterative optimization, decision-makers need to make their way through many back-and-forth iterations until they reach a desired system design. As the components are visited one after the other, decision-makers also need to think multiple steps ahead to anticipate the implications of their current choice in the following iterations. As a consequence, analysts might tend to proceed with the first working solution rather than striving to find better designs [44]. A simultaneous exploration across individually modeled but dependent components is needed to make system design more efficient.

With a system, the decision task turns into a series of multi-attribute choices, one for each component involved. Each of these choices might have side effects on the system's operability and performance.

2.4.3 *Engineering Design Properties*

Since the late 1990s, decisions have been increasingly recognized as a fundamental construct at the core of the engineering design process [93]. Ullman even claims that *"design is decision-making"* [281]. Different models have been proposed to describe the engineering design process. Commonly involved steps are to 1) identify the problem, 2) gather information, 3) generate alternative solutions, 4) evaluate the alternatives and decide, and 5) communicate the results [62]. Steps one, three, and four correspond to the stages of intelligence, design, and choice according to Simon's model (Section 2.1.3). Some works also reduce the process to the steps of alternative generation and alternative selection [93]. As decisions in engineering design are targeted at real-world objects, they should be directly applicable to the domain's practice. They are also of high consequence, i.e., they may have critical impacts on their surroundings. As such, engineering design requires a careful management of data, models, and knowledge to inform a judgment, on which the decision is based.

Whether targeted at a system, component, or process, engineering design can be characterized by four properties [108]. First, it is *purposeful*, i.e., guided by an explicit goal. This refers to some functions or properties that the subject to be designed should perform or have. Second, designs are shaped by *specifications* and complex *constraints* that incorporate a lot of engineering domain knowledge. Specifications tell what the designed subject is intended to accomplish, while constraints are limitations that the decision-maker must stay within. Some are absolute hard constraints, but many are relative and must be balanced against each other and against how well a design satisfies the specifications. Third, the process of engineering design is *systematic and iterative*. Although there are no imposed instructions, we can observe recurring characteristic steps, which are repeated as necessary (e.g., the steps listed above). Otherwise, decision-makers would most probably get lost in an endless search for the optimal solution (compare analysis paralysis in Chapter 1). Finally, engineering design is a *social and collaborative* undertaking, e.g., as commissioned work (an example case is detailed in Chapter 6) or as group decisions. In any case, interdisciplinary work and communication with stakeholders are crucial. In the course of this thesis, it will become apparent that these properties largely hold true also for multi-attribute choice in general.

2.4.4 *Potential for Visualization Support*

The characteristics described above clearly indicate the potential of using visualizations to support engineering design. Although (evolutionary) multi-objective optimization algorithms are commonly used, Hazelrigg claims that engineering design "*can never be reduced to a*

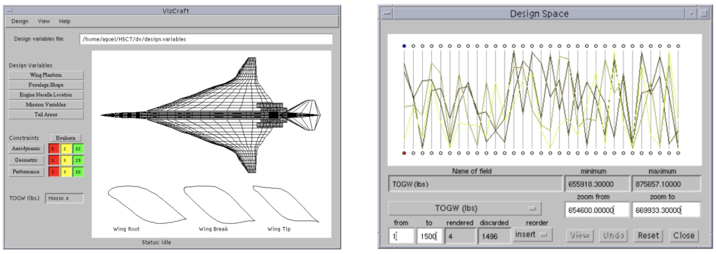

(a) VizCraft, 1999 [80]

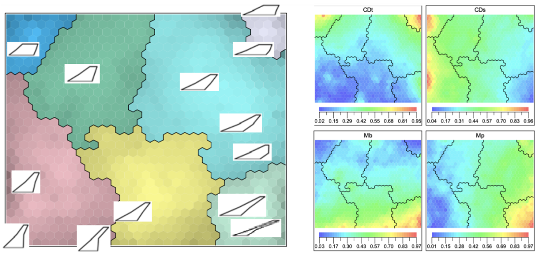

(b) Obayashi et al., 2003 [192]

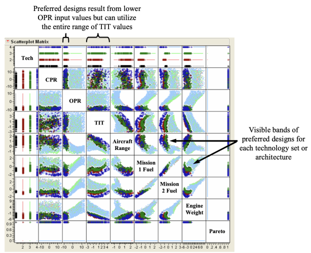

(c) Marsaw et al., 2007 [160]

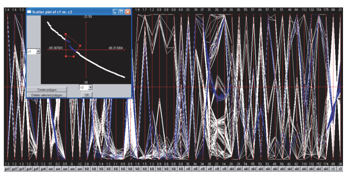

(d) Kipouros et al., 2013 [126]

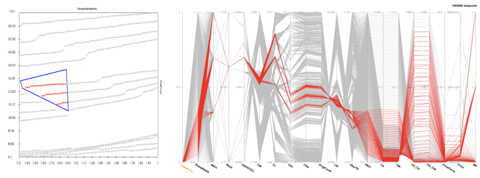

(e) Piotrowski et al., 2019 [208]

Figure 2.7: Examples of visualizations for design in aerospace engineering.

prescriptive procedure exclusive of human input" [93]. Human input can include relevant relationships to be considered, an appropriate goodness metric, or judgment regarding which options to consider in detail. According to Hazelrigg, decisions in engineering design should be based on information obtained from a variety of sources [93]. Visualization can make these information accessible to decision-makers and serve as an interface to human input. It can help make informed decisions by highlighting relationships and trade-offs, conveying their nature, and making the effects of applied constraints and preferences visible. The value of visualization has been recognized in engineering disciplines, where researchers have come up with different interactive visualization solutions themselves. Even if these solutions do not originate from the field of visualization research, they can be a valuable source to inform our problem-oriented methodology.

Early works published at the IEEE Visualization conferences in the late 1990s offer a compelling mixture of practical experience and trend-setting outlooks regarding the role of visualization in engineering design. Spence et al. highlight the absence of a direct mapping from design performance to design parameters [255]. They propose the *Influence Explorer*, which overloads interactive histograms on parallel attribute axes to reveal the relationship between the design and objective space. Brushing and linking helps solve design-specific tasks such as inspecting individual options, revealing harmony or conflicts among attributes, specifying performance requirements, and considering manufacturing tolerances. Shaffer et al. highlight the impracticability of a complete search in the design space of an aircraft design problem [244]. They claim that visualization can serve multiple purposes in engineering design: provide insights into a design problem by provoking new interpretations, provide a means to effectively search the design space, and provide a means for engineers to apply their design expertise to a problem. They also propose a number of requirements for a visualization system to support aircraft design. These requirements might also generalize to other engineering design problems. Based on these initial thoughts, Goel et al. propose *VizCraft* to optimize aircraft designs towards minimal take-off gross weight [80]. It combines parallel coordinates for an overview of the design space with a detail view of the geometry, design constraints, and performance indicator of a selected option (Figure 2.7a).

Design spaces in aerospace engineering are often spanned by tens or hundreds of (geometric) design parameters. One way to reduce this dimensionality for analysis is a projection to two-dimensional visual space, for example using a self-organizing map (SOM) [132]. Obayashi and Sasaki use a SOM to project a 72-dimensional design space of supersonic aircraft wings to two visual dimensions, resulting in seven clusters which are overlayed with their representative wing shapes [192] (Figure 2.7b). A side-by-side comparison of this SOM color-coded by each of the four criteria allows engineers to investigate the criteria-wise performance of wing shapes but also to identify relations among criteria in certain regions of the design space. Marsaw et al. decided for a lossless projection using a scatter plot matrix to study three design parameters of an aircraft engine with respect to five criteria [160] (Figure 2.7c). They target an early design phase, where a long-term decision about the basic engine architecture for future parametric studies is made. Color-coding of the points in the scatter plot matrix is used to distinguish four possible architectures. Kipouros et al. employ parallel coordinates for a lossless overview of the relation between tens of design parameters and two to three criteria [126] (Figure 2.7d). They particularly appreciate this visualization technique for its manifold interaction possibilities that facilitate to center engineering design processes around the human decision-maker. Range selection is used to identify design parameter values leading to desired performances. Piotrowski et al. extend the interactive capabilities of parallel coordinates by introducing proportion-preserving simultaneous manipulation of multiple range brushes or superimposed composite brushes to represent different sets of constraints [208]. They also link the parallel coordinates to scatter plots of any pair of axes, where free-form polygonal brushes can be specified (Figure 2.7e).

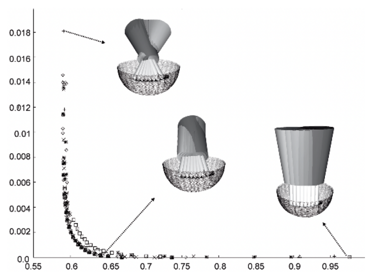

(a) Müller et al., 2008 [183]

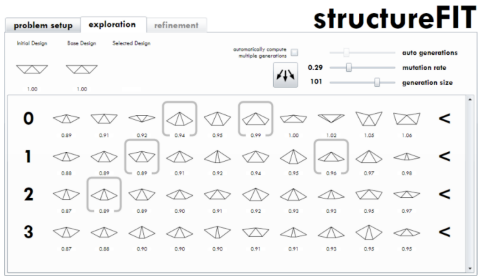

(b) structureFIT, 2015 [180]

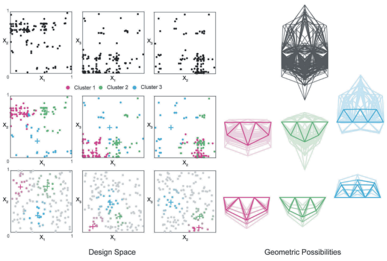

(c) Brown and Mueller, 2017 [40]

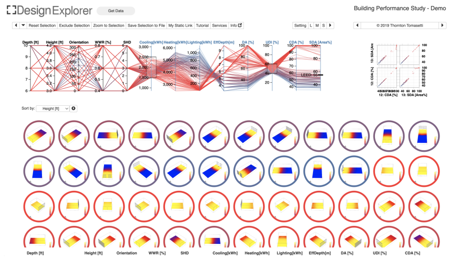

(d) Design Explorer, 2019 [312]

Figure 2.8: Examples of visualizations that depict quantitative attributes alongside qualitative aspects like shape or geometry.

Composite brushes across visualizations either represent intersections or unions of the individual brushes.

In engineering design, quantitative attributes often need to be considered alongside qualitative aspects like shape or geometry. Many tools therefore employ multiple views that allow decision-makers to simultaneously view the solution space from different perspectives: the abstract design space, the resulting three-dimensional geometry, and the abstract objective space. Müller et al. highlight the importance of a joint exploration of these spaces for gaining insights into the characteristics of an optimization problem and choosing a final solution [183]. They showcase the application of well-established visualizations like parallel coordinates or scatter plots on the optimization of a five-axis milling process (Figure 2.8a) and the design of cooling layouts for molding. In the domain of electric motor design, most commercial tools provide only two-dimensional Pareto front visualizations that are not suited for optimization with multiple attributes. An exception is *SyMSpace*, where motor designs are explored in the objective space in a scatterplot matrix that conveys all pairs of attributes to be optimized [248]. Relevant options can be brushed and observed in the design space using linked histograms. Still, the analysis in SyMSpace is limited to two-dimensional projections of the Pareto front.

While optimization approaches have settled in fields like aerospace and mechanical engineering, computationally methods also increasingly enhance design processes in more creative fields like architecture. *structureFIT* depicts the generation-wise evolution of candidate designs for a truss roof by arranging their structural geometries together with a unidimensional performance indicator in a matrix [180]. By selecting one or more designs in the overview and adapting the mutation rate and generation size, the decision-maker can steer an optimization algorithm towards higher performing versions of a selected design or more distant regions of the design space (Figure 2.8b). Brown and Mueller survey ambitions towards data-driven architectural design, where computational methods and human decision-makers tightly collaborate to explore the relation between design and objective space [40]. These include scatter plots to visualize clusters of structurally similar designs (Figure 2.8c) or small multiples of parallel coordinates to visualize selected design parameters of similarly performing options. An increasing interest in the field of data visualization has given rise to dashboard tools like the *Design Explorer*, which links views of the structural geometries to parallel coordinates and scatter plots depicting the abstract design and objective spaces [312] (Figure 2.8d).

The visualizations emerging out of the field of engineering design offer important insights into the way engineers think about decision-making in design. Still, as they are not visualization researchers, engineers might not be aware of the breadth of approaches proposed in visualization literature. However, a solid knowledge of the visualization literature could significantly help "*broaden the consideration space of possible solutions, and [to] select good solutions over bad ones*" [243]. Additional potential might thus be unlocked by visualization and engineering researchers joining forces. However, while researchers in engineering design frequently employ visualizations to make choices, problem-driven research in the visualization domain has not yet been dedicated to this field.

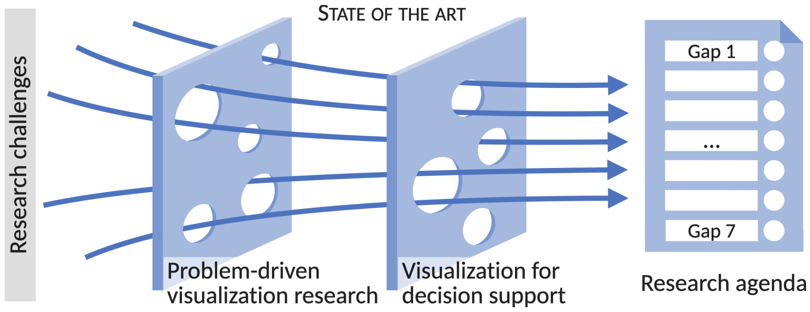

Figure 3.1: This chapter derives research gaps that remain from the research challenges (Section 1.2) after reviewing prior works on design study methodologies and visualizations for decision support.

3 STATE OF THE ART IN INFORMATION VISUALIZATION

THIS chapter provides the scientific foundation in information visualization, forming the basis for the research presented in this thesis. It complements the theoretical principles of multi-attribute choice tasks introduced in Chapter 2 by providing methodological and technical context necessary to research visualization-based choices.

What the reader can expect from this chapter:

- A literature review of design study methodologies revealing that guidance on how to elicit expert knowledge and assess long-term usefulness remains vague (Sec. 3.1).
- A literature review on multivariate visualizations with varying decision support claims showing that although they support basic analytic and decision-making activities, they do not meet the information needs of expert decision-makers yet (Sec. 3.2).
- A research agenda comprising seven research gaps regarding decision support visualizations (Sec. 3.3).

3.1 PROBLEM-DRIVEN VISUALIZATION RESEARCH

Aiming to design visualizations for solving real-world choice problems means to conduct problem-driven visualization research, most commonly represented by design studies. A *design study* is a project where visualization researchers analyze a real-world problem in a target domain, design a validated visualization solution for it, and reflect on lessons learned to refine guidelines [243]. For this, researchers engage closely with data expert users who have deep knowledge in the target domain [296].

3.1.1 *Design Study Methodology*

Since the first call for design study papers [187], we have seen an increasing interest in systematic problem-driven visualization research. A first definition of the building blocks of a design study has been proposed with the four levels in Munzner's *nested model*, which describe a tight coupling between the creation of visualization designs and their evaluation [185]. The nested model was later refined by dividing each level into descriptive blocks and introducing guidelines that connect blocks within and across levels [171]. Marai expanded the domain characterization stage of the nested model by functional specifications [158]. The *design activity framework* links four design activities (understand, ideate, make, and deploy) to the levels of the nested model [164]. Sedlmair et al. provide methodological and practical guidance on how to systematically conduct a design study. They propose the *nine-stage framework* to describe the entire process ranging from selecting promising collaborations to publishing a design study paper [243]. It was adapted by Syeda et al. to make the process fit the duration of one semester for use in visualization pedagogy [261]. Meyer and Dykes provide guidelines to ensure rigor in conducting design studies [168]. Arbesser et al. provide a high-level view on challenges, strategies, and lessons learned from distributing visualization tools to domain experts in different application areas [13].

A design study typically builds upon existing visualization design models and evaluation methods. Regarding visualization design, Lloyd and Dykes report on long-term experience with human-centered design methods [151]. Tory and Möller review human factors research in the context of visualization design [270]. Other works focus on the interplay between data, users, and tasks [176] or gaps and knowledge precepts underlying visualization design [7]. A variety of approaches has also been established for evaluation. Carpendale provides an overview of relevant methods to evaluate information visualizations [45], while Isenberg et al. review their actual use in visualization research [113]. Lam et al. describe the use of evaluation methods for different purposes in seven guiding scenarios and identify realism in tasks, data, and users as critical for validity [138]. Aiming to evaluate visualization designs in open-ended real-world settings generally suggests the use of qualitative methods, such as case studies [246], insight-based evaluations [231], grounded evaluation [112], or expert reviews [271]. Contributions achieved through a design study can take various forms, from learnings about the problem domain (see Section 3.1.2) over validated visualization designs to a refinement of methodological guidelines [168, 236]. An overarching goal of design studies is generalizability, i.e., the potential of the outcome being transferred to other contexts. Generating these forms of contributions from design practice is commonly achieved through the process of reflection [167].

3.1.2 *Domain Characterization Methods*

Section 3.1.1 provides an overview of design study methodologies and implementations thereof in the context of visualization-based decision-making. This section details the very first step of a design study, the *domain characterization*. It is also referred to as *context of use* [107], *task elicitation* [194], or *domain problem characterization* [185]. In this stage, visualization researchers aim to understand the data, tasks, goals, and experiences of the domain experts as well as the conditions and constraints that will frame the visualization use [243]. Domain characterization is conducted *before* the design stage. Its outcome is often implicitly defined and includes data-user-task definitions [176], dominant concepts as a result of coding [272], or design implications [138]. Translated to domain-agnostic abstractions, they guide the visual encoding decisions in the subsequent steps of the design study [185].

Such guidelines depend on the quality of the domain characterization *activities* performed by visualization researchers. Thus, if researchers want to design useful problem-driven visualizations, they need to know how to conduct domain characterization effectively. Visualization researchers typically read about the domain [54, 151, 169, 239], gain experience with domain tools and data themselves [23, 166, 242], or perform a variation of "talking with and observing domain experts" [243]. Engaging with domain experts includes interviews and real-world observations [29, 139], contextual inquiries [151, 239, 241], focus groups [239, 241], or workshops [82]. To understand more concretely what those activities entail, we reviewed around 30 papers on problem-driven visualization research (including the design studies reviewed by Sedlmair et al. [243]).

Most domain characterization reports focus on the outcome of the activities, i.e., data and task abstractions, rather than how they arrived there. Descriptions of the process remain fairly high-level. They range from "interviews with application experts" [209] over having "interviewed experienced assay developers" [215] to having observed "real model developers" [29], "the domain expert on a real-world use case" [54], or "daily work practices" [237]. Some reports are even limited to having collaborated with domain experts over a certain time period [35, 170, 172]. Reports sometimes include the type of interviews (e.g., guided [242], semi-structured [139, 239], or unstructured [282]) as well as the researchers' topics of interest (e.g., workflows [20, 54], data and analysis methods [169], or tools [166, 242]). Few works explain the interview contents, goals, and subsequent data analysis steps in detail [68, 139, 151, 241]. Few also provide rich reports of the observation procedures they followed and the qualitative insights they gained [239, 241, 242]. Still, it often remains unclear whether these descriptions reflect the individual choices of the researchers in that specific context or an established methodological protocol, which could be re-used beyond the specific design study.

This limited methodological justification, or under-reporting, of the domain characterization stage can be attributed to the lack of structured guidance on how to study an application domain in the context of visualization design studies. Munzner observed that "hardly any papers devoted solely to analysis at this level [domain characterization] have been published in venues explicitly devoted to visualization" [185]. Marai states that "although visualization design models exist [...], these models do not present a clear methodological prescription for the domain characterization step" [158]. She proposed an actionable framework for domain characterization, which centers around activities and tasks. Yet, we still lack explicit guidance on how to extract the expertise and experience involved with problem-solving.

Current practices are grounded on the diversity of methodologies used to study people, cultures, and habits in ethnography [185]. However, these methodologies have not been developed against the background of data analysis. To convey the intended message with a visualization, we need to trace how an expert applies domain knowledge to interpret the depicted information [48]. This understanding is difficult to obtain. Domain expertise and experience rely heavily on personal *tacit knowledge* [212, 287], which involves contextual implications, analogies, or judgments of typicality. Unlike explicit knowledge that has been verbalized, written down, or stored in a database, tacit knowledge is hidden in users' minds. It cannot be derived from observable behavior and users find it hard to articulate how they do something that is based on expertise [212]. Thus, tacit knowledge can only be acquired from humans through their cognitive processes [75]. Despite recent advances in ethnographic methodologies, task taxonomies [38], and analytic question sets [138], we identify a lack of prescriptive steps for visualization researchers to follow in a design study in order to elicit domain knowledge and derive task abstractions.

3.1.3 *Long-Term Evaluation Methods*

As design studies aim to solve a real-world problem in a target domain, one form of success is the long-term adoption of the designed visualization. The domain threat in problem-driven visualization research is that users do not in fact have the targeted problem, while the abstraction threat is that the chosen generic data types and analysis tasks do not solve the identified problem [185]. As a downstream validation to both threats, Munzner proposes to observe the adoption rate and document long-term use of the deployed visualization design [185]. This helps to assess factors that influence how a tool is adopted in the intended work environment [138]. For example, it helps to find out whether the barrier to regular day-to-day usage is an integration issue, or an indication that the tool failed to address the true needs of the target users [37]. If domain experts repeatedly and on their own initiative use the tool in their daily work, this strongly suggests that they indeed face the targeted problem and that it lastingly benefits from the visualization support.

Longitudinal studies also account for the fact that tool usage might be different between the summative evaluation at the end of a design study and the expert's day-to-day work that follows the design study (publication). A summative evaluation often implicitly or explicitly asks domain experts to use the tool as opposed to the self-initiated use at the experts' workplaces whenever they consider it beneficial [185]. As with other newly introduced tools, it takes a certain time for target users to adjust to a visualization support, such that it is operated in their workflow under stable and normal conditions [91]. Collecting anecdotal or empirical evidence of a tool's short-term usefulness might not generate insights about its routine use and long-lasting limitations or benefits. For example, this might be the case for the reflection on benefits, challenges, and potential improvements of *WeightLifter* one month after deployment [198]. Anecdotic evidence suggests that, despite promising initial acceptance, visualization use typically decreases after a while [37] – most design studies do not follow up upon that.

An exception is the work by Gonzales and Kobsa, who conducted a longitudinal study to clarify the true adoption of their tool after preceding empirical studies had shown promising results [81]. They found that domain experts gradually lost interest in integrating the tool into their current work routines, which they attributed to a misunderstanding of the experts' workflow as a potential factor. Similarly, Kang and Stasko complemented an earlier lab study by interviewing analysts with different real-world tasks who they knew had used their tool on their own initiative for two to 14 months [121]. Studying the adoption of *Overview* effectively helped Brehmer et al. refine their understanding of why and how domain experts used their tool, ultimately leading to revised task abstractions and design rationales [37]. Their design study involved multiple deployments and analyses of user adoption over a period of two years. Kincaid et al. collected user feedback from different research labs that had employed their visualization tool for scientific studies, one of them having used the tool for more than a year [125]. McKeon observed the activities of general web users on a public deployment of their wiki-like visualization dashboard system over half a year (including a beta phase) [165].

Longitudinal field studies have also been conducted without a particular focus on adoption. Shneiderman and Plaisant advocate for measuring the utility of a visualization by the success it helps domain experts achieve in carrying out leading work in their fields [246]. To this end, they propose guidelines to evaluate the prolonged use of information visualizations over several weeks or months. In line with their guidelines, many longitudinal studies collect diverse data, from audio and notes over screen captured interaction sessions to logging data. They have studied insight generation [231], task abstractions [166], the integration of statistics with visualization for exploratory data analysis [205], and early stages of geovisualization design [151].

The nature of longitudinal studies fosters incremental tool improvements and learnings. Still, except for the work by Brehmer et al. [37], studies spend little effort on reflecting how feedback evolved across different post-deployment stages and how this might relate to changes made to the tool, e.g., whether a change has actually solved a problem. Most post-deployment studies also suffer from survivorship bias (a particular type of selection bias), i.e., the studies solely focus on participants who had used the tool [37, 125, 165]. Instead, we can also learn from following up on unsuccessful cases, in which experts at some point lost interest in using a tool.

To summarize, visualization tools are typically evaluated with respect to their short-term usage only. Although a few works on long-term post-deployment evaluation exist, they are still remarkably rare [37, 185, 246] given that a core goal of design studies is to develop lasting solutions to problems in a target domain. Although an appropriate timing is critical [91], existing longitudinal studies rarely quantify the evaluation period or the tool usage period preceding the evaluation. Still framing their research as long-term studies, this makes it impossible to know what time period the authors refer to. They also do not yet exploit the full potential of learning from why adoption failed and how tool usage evolved over multiple post-deployment stages.

3.2 VISUALIZATION FOR DECISION SUPPORT

Without prior knowledge, the closest we can get to solving a multi-attribute choice problem is to restrict the choice to incomparably good alternatives. This is typically the Pareto front, i.e., those alternatives that cannot be improved in one attribute without impairing at least one other attribute (Section 2.2.2). In most cases, the information can be structured in the form of a multivariate data table, where the rows are alternatives and the columns represent the attributes. The final decision depends on the decision-makers' interpretation of the data to determine the most preferred alternative. The interpretation of the data, in turn, heavily depends on their subjective prior knowledge.

The number of alternatives to be considered can be very large. To support a decision-maker in finding the most preferred compromise, the alternatives can be visualized [175]. Korhonen and Wallenius differentiate between visualizing a single alternative, a finite set of alternatives, or an infinite set of alternatives [135]. Because the number of decision criteria is often lower than that of design parameters, special attention is usually paid to visualizing the alternatives' performances [173]. Visual aids can also convey the relative importance of attributes, the relationship between attribute weights and alternative utility, and immediate changes of the above to observe the dynamics of the decision process [134]. Providing a holistic view of alternatives requires visualization methods that effectively map large multivariate data to two-dimensional visual space [152]. Their design might be informed by findings from the cross-domain review conducted by Padilla and colleagues [197]. They studied the role of working memory in decision-making with (static) two-dimensional visualizations and how it interacted with fast, light-weight decision-making versus contemplative, effortful decision-making.

In this section, we review interactive visualizations that are relevant in the context of decision-making. Following the data-user-task design triangle [176], the review covers the following scope.

Data-wise, we focus on decision problems where the set of multi-attribute options is finite and known in advance. Interactive optimization approaches where decision-makers steer the refinement of solution sets are investigated in a different research area [88, 89]. The targeted decision problems are largely based on multivariate or multi-run data. While the data might be spatio-temporal in that it contains information about space or time associated with an option, we exclude dynamic data that changes over time [4]. We consider the attribute values of options certain. Uncertainty associated with the data largely affects decision-making and requires special consideration in visualization support [28, 34]. We further exclude multi-modal data, where the same subject of interest is sampled at different temporal or spatial resolutions, and multi-model data, where different but related subjects of interest are sampled or computed [122].

Task-wise, we focus on multi-attribute choice tasks, i.e., the task of selecting the best among a number of options. We thus exclude works that target only the intelligence and/or design stage of the decision-making process (see Section 2.1.3). Visualization tools supporting these stages have been covered elsewhere [195]. We also exclude works that investigate the dynamics of evolutionary processes leading towards a solution set [105].

User-wise, we focus on single decision-makers as opposed to groups of decision-makers. Group decision-making raises particular demands regarding consensus-building, which has inspired the design of dedicated visualization tools [17, 99]. Section 2.3 has pointed out the differences between single target users being analysts, casual decision-makers, or expert decision-makers. Following these categories, we distinguish between general multivariate visualization approaches without a claim to support decision-making (Section 3.2.1), visualizations that have been illustrated with common multi-attribute choice tasks like buying a car (Section 3.2.2), and visualizations explicitly designed to support expert decisions (Section 3.2.3).

3.2.1 *Multivariate Visualizations*

As explained above, the finite set of alternatives underlying a multi-attribute choice can be considered a multivariate data set. Here, we review general information visualization approaches that are meant to help understand this type of data set. As most of them have not been exclusively developed against the background of decision tasks, we also discuss their relation to multi-attribute choice. Multivariate visualization approaches related to decision-making can be grouped into dimension reduction techniques and lossless geometric projection [135]. We exclude non-geometric approaches, e.g., icon-based or pixel-oriented visualization techniques, for their disadvantages pointed out by Dimara et al. [64]. We also summarize how multivariate visualizations for parameter space analysis can help decision-making.

Multivariate Visualizations Using Dimensionality Reduction

Dimension reductions provide a dense representation of virtually any number of dimensions. A popular approach to dimension reduction is the self-organizing map (SOM) [132]. In engineering design, the SOM has been employed for criteria-wise design space exploration in the context of aerodynamic optimization [191]. 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 [49]. This is similar to the *RadViz* approach, where the mapping of alternatives is defined by attribute anchors acting as springs that exert forces on the alternatives [100]. In *Dust & Magnet*, the attribute anchors can be freely positioned and act as magnets, where alternatives with high attribute values are attracted faster to a magnet than alternatives with lower attribute values [304]. The *Data Context Map* also contains arbitrarily positioned attribute anchors and overlays semi-transparent regions in the projection that satisfy different (desired) attribute ranges [50]. Zhao et al. use t-SNE as a projection method for their system *SkyLens* and, like SOMMOS, employ radial glyphs to represent individual alternatives [309]. These tools will be revisited in the context of decision scenarios in a later section. While dimension reduction is useful for navigating Pareto fronts, the resulting dimensions and visual layout are often hard to interpret [238] and raw attribute values of alternatives cannot be read directly. However, the personal judgment and choice of alternatives generally requires users to be able to retrieve attribute values by looking at the visualization. Dimension reduction also helps perceive groups of similar or dissimilar alternatives. However, having filtered the alternatives into acceptable regions, decision-makers are often more interested in a direct comparison and trade-off between alternatives. For these reasons, we do not primarily consider dimension reduction approaches for multi-attribute choice.

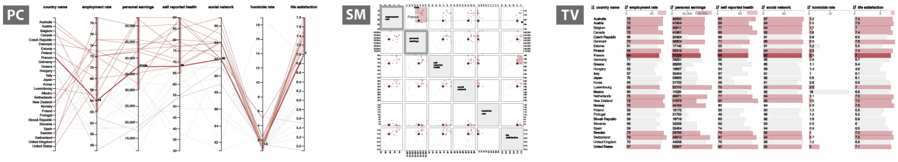

Figure 3.2: Parallel coordinates (left), scatterplot matrices (center), and tabular visualizations (right) are common visualization techniques that allow decision-makers to retrieve any attribute value of any alternative without interaction. Image from Dimara et al. [64].

Multivariate Visualizations Using Lossless Projection

In contrast to dimension reduction, a lossless projection enables decision-makers to visually retrieve any attribute value of any alternative without interaction [64]. This category includes common multivariate visualizations that use rather simple visual encodings to represent raw attribute values in a single plot. One example are scatterplots. Mühlbacher et al. use scatterplots to visualize the trade-off between accuracy and complexity of decision trees [181]. For extension to multivariate data, scatterplots for every pair of attributes can be arranged in a scatterplot matrix (Figure 3.2, center). Scatterplot matrices have been employed, e.g., for the design of an aircraft engine [160]. Elmqvist et al. propose animated rotations in 3D space to navigate through the multi-dimensional space by transitioning between the two-dimensional projections of neighboring scatterplots [73]. Scatterplot matrices effectively convey correlations and comparisons between selected attributes. However, their complexity rapidly grows with increasing number of attributes and the perception of multi-attribute alternatives is limited to pair-wise projections.

Another example are tabular visualizations. Cell values are encoded visually, with a popular encoding being length (i.e., bars) [206]. The resulting layout is related to stacked bar charts, except that bars are aligned to a separate baseline for each attribute rather than stacked. While some tools allow to switch between these layout strategies [84], stacked bars make it difficult to compare individual attribute values across multiple options. Tables with aligned bars have been found to perform well for multi-attribute choice tasks in terms of completion time and technique preference [64] (Figure 3.2, right). Cells that are shaded by hue or luminous intensity make the tabular visualization a heatmap [216]. Tabular visualizations provide an overview of multiple alternatives and attributes simultaneously. Dedicated table layouts also exist for hierarchical data [145], although this might remain a boundary case in multi-attribute choice.

The predominant technique to visualize Pareto fronts are parallel coordinates [16]. Attributes are mapped to parallel axes and alternatives are represented by polylines that intersect the axes according to the respective attribute values [109] (Figure 3.2, left). Parallel coordinates are particularly suited to convey the characteristics of attributes and the nature of their relation, e.g., "it seems that low power loss inevitably means high price". Andrienko and Andrienko propose axis modifications regarding orientation, scaling, alignment, and ranking of alternatives [10]. Other variations propose different layouts of axes (e.g., circular, hierarchical) or polylines (e.g., curved or bundled) [117]. Parallel coordinates present a compact two-dimensional visual representation that allows for a perception of alternatives across all design parameters and decision criteria. They are also a popular choice by domain experts, e.g., in engineering, who are not necessarily visualization researchers or designers. Among others, parallel coordinates have been applied to decision problems in automotive engineering [28, 163], aerospace engineering [80, 244], and aerodynamic engineering [126], also in virtual reality settings [262]. Fleming et al. reflect on the use of parallel coordinates from the perspective of multi-criteria optimization in real-world engineering design [77]. The horizontal line up of vertical axes in parallel coordinates allows for many attributes being shown next to each other. In addition, parallel coordinates particularly expose conflicting attributes in the form of crossing line sections between the adjacent axes.

Once a few promising alternatives have been identified, star plots are often used for their detailed comparison [309]. Their visual encoding is equivalent to that of parallel coordinates except that the axes are arranged radially. Unlike parallel coordinates, which are suited to identify relations between attributes, star plots are better suited for comparing specific options [256]. Star plots have also been used to depict alternatives resulting from an evolution of the same system over time [207]. To include the relative importance of attributes in the display, the axes of a star plot can also be shortened and stretched according to weights [263]. A similar strategy has also been proposed for radial bar charts, where the angular size of a bar can be proportional to the weight of the respective attribute [263]. Due to the compact circular arrangement of axes and segments in a radial layout, the number of attributes that can be effectively perceived is limited.

Visual Parameter Space Analysis

Section 2.4.1 pointed out that a direct mapping from a desired performance to a design option does not exist. Gaining insight into the correspondence between design options and performance is thus an important aspect of multi-attribute choice [244]. How does choosing one design option over another influence the performance? Which performances are generated by which kinds of design options? Similar questions are addressed with parameter space analyses [240]. These analyses are motivated by the use of input-output models, which map a set of input parameters to a set of outputs, most often resulting in multirun data [122]. One goal of parameter space analysis might be model validation, e.g., to determine how well a regression-based model approximates the output of a more detailed simulation [182,

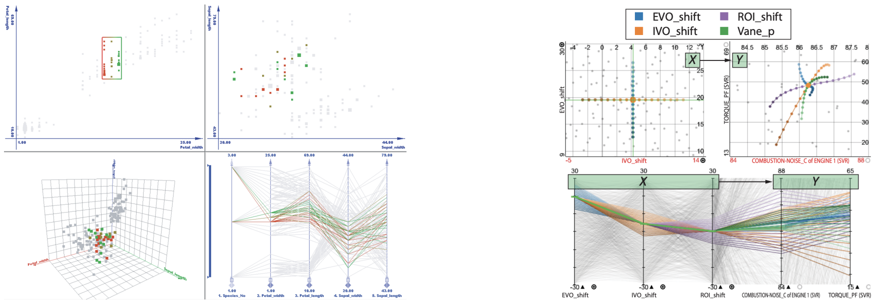

(a) Matković et al., 2005 [163]

(b) Berger et al., 2011 [28]

Figure 3.3: Examples of tools that help explore the relation between input parameter settings and the corresponding output behavior.

209]. Another goal that relates more closely to multi-attribute choice is to choose an input parameter setting that results in a preferred output. This typically requires an understanding of the relationship between input parameters and outputs in both directions. Derived from a meta analysis of 21 design studies, Sedlmair et al. provide a conceptual framework for parameter space analysis that includes a data flow model, navigation strategies, and analysis tasks [240].

The current body of work in visual exploration of parameter spaces mainly comprises problem-driven investigation of complex systems in different application domains like meteorology [213] or parametrized image analysis [215, 268]. The *Influence Explorer* showcases an early overloading of histograms on parallel coordinates to support interactive exploration of computed performance values for different variants of a product in engineering design [279]. Another early work has been presented by Shaffer et al., who accurately characterize the role of visualization in analyzing the high-dimensional parameter space related to aircraft design [244]. Some of the characteristics expressed in their work were taken up in later works related to parameter space analysis. Matković et al. link multiple basic information visualization techniques for exploring the parameter spaces of fuel injection systems [163] (Figure 3.3a). Informed by three distinct application domains, *ParaGlide* proposes to interactively divide the input parameter space into partitions that represent distinct output behaviors [29]. In this way, users can identify input regions producing desirable outputs as a starting point for detailed trade-off analyses.

Parameter space analysis aims at an option with a set of design parameter values that meet desired performance requirements. Where this option does not exist physically in the real world yet (e.g., choosing a prototype to be built as opposed to choosing from existing universities), it is constructed from these design parameter values. In this case, decision-makers need to be aware of tolerances. They introduce uncertainty because they permit small deviations between the nominal design parameter values and the product's actual properties after construction. How such a variation might affect the product's functioning needs to be kept in mind already during decision-making. Studying how the performance of an option changes in response to perturbations or uncertainty in its design parameter setting is called sensitivity analysis [230]. For multi-attribute choice, approaches where derivations in the design parameters are investigated within a small neighborhood of the chosen option are most relevant. At the intersection between sensitivity analysis and decision-making, Berger et al. enable users to navigate a continuous multivariate design space towards preferred parameter regions with respect to multiple criteria [28]. Originating from a focal design, star sampling in the design space is performed to observe criteria ranges that are within reach, eventually hinting at potential improvements or necessary trade-offs (Figure 3.3b). 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. Focusing on the neighborhood of a focal design for sensitivity analysis, this approach does not provide a comprehensive overview of the entire parameter space.

Parameter space analysis and sensitivity analysis play an important role in informing a multi-attribute choice. However, many approaches build upon interactive steering of the parameter space sampling during analysis. This is rarely applicable in multi-attribute choices, which aim at deciding among a predefined set of options. Furthermore, an understanding of the relationship between inputs and outputs alone is not sufficient to make a final choice. The patterns, hypotheses, and insights generated through analyses need to be viewed against the decision-maker's preferences to become actionable and to move from a desired parameter region to a final decision. While research around parameter space analysis and sensitivity analysis can be a valuable source of inspiration and lessons learned, the approaches are not directly transferable to the task of making a multi-attribute choice.

Summary

Given the large body of visualization techniques for multivariate data, not all of them are useful for visualizing alternatives in the context of decision-making [175]. Well-designed tools for analytic tasks like identifying outliers or patterns in multivariate data can serve as an important foundation for decision support. But their conception might not take into account the psychological effects or cognitive biases that might affect the final choice in an irrational way [60, 197]. A solid knowledge of general multivariate visualization techniques is important for researchers targeting decision tasks. Still, the diverse subjective prerequisites that take effect in choice situations require decision support tools that result from dedicated considerations regarding visual encoding, navigation, and filter strategies.

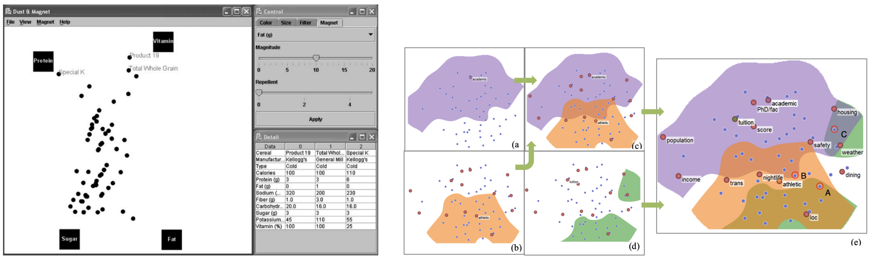

(a) Dust & Magnet, 2005 [304] (b) Data Context Map, 2015 [50]

Figure 3.4: General-purpose visualization tools for casual decision-makers that make use of dimension reduction remain an exception.

3.2.2 *Casual Decision-Makers: General-Purpose Tools*

From the general techniques of visualizing multivariate data, we move to visualization tools specifically meant to support multi-attribute choices made by non-experts. These are personal decisions that many of us face at some points in our lives. For example, as consumers, we often need to decide for a product from an online catalog, such as which film to watch or which camera to purchase. In contrast to such micro-decisions, personal decisions can also take the form of permanent life decisions, such as which house to buy or which university to attend.

Visualizations for a general audience of non-experts need to be simple and clear, with reduced complexity regarding visual elements and interactions. We organize our review of casual visualizations for common choice tasks roughly along the lines of the previous review of general multivariate visualizations, i.e., considering dimensionality reduction and lossless projection but also approaches to utility scoring.

Casual Decision Tools Using Dimensionality Reduction

We will soon see that visualizations supporting common multi-attribute choice tasks typically employ lossless projections. One exception to this strategy is *Dust & Magnet* [304]. Multivariate options are projected to 2D space and magnet anchors are placed and moved in the same space to separate desired from undesired options, e.g., to choose cereals based on their dietary composition. If magnets for protein and vitamin are placed far away from magnets for sugar and fat, they will attract healthy cereals with high protein and vitamin and low sugar and fat (Figure 3.4a). Another exception is the *Data Context Map*, which additionally visualizes regions of desired attribute ranges in the projection [50]. This allows decision-makers to identify overlaps where, e.g., universities or cars provide the best trade-off among preferences regarding multiple attributes (Figure 3.4b). Still, these representations are lossy projections, such that raw attribute values of options can only be retrieved via details-on-demand.

Table 3.1: The reviewed lossless projection visualizations for decision support of casual users organized by year and primary visualization. Tools marked with 🏅 build upon utility scoring as opposed to incremental querying. Arrows indicate tools inspiring one another.

Year	Tabular Visualization	Scatter Plot	Parallel Axes
2020			ConfigurationFinder [222]
2018	🏅 PODIUM [284]	🏅 ReACH [295]	TOP-Slider [142]
2013	🏅 LineUp [85]		
2012	ManyLists [149]		Product Explorer [221]
2008	🏅 VDM [302]	ScatterDice [73]	
2007		🏅 VMAP [266]	
2006	🏅 ValueCharts+ [24]		
2004	🏅 ValueCharts [43]		
2001			EZChooser [298]
2000		SmartClient [218]	MultiNav [141]
1996	FOCUS [257]		
1994		FilmFinder [2]	Attribute Explorer [280]
1992		HomeFinder [297]	

Casual Decision Tools Using Lossless Projection

In contrast, retrieving raw product properties from a catalog represented as a lossless projection, e.g., based on scatter plots, tabular visualizations, or parallel coordinates, does not require interaction. The reviewed lossless projection visualizations targeting casual decision-makers are summarized in Table 3.1.

An early scatterplot-based approach that has inspired many works is *HomeFinder* [297]. Hundreds of homes for sale are represented as dots on a 2D geographical map and can be filtered based on attributes using sliders in a control panel (Figure 3.5a). Reachability can be considered by marking locations on the map and defining queries on the linear distances to these points of interest. The reachability estimation has later been refined in *ReACH* (Figure 3.5b) by considering traffic conditions as well as more fine-grained constraints representing daily routines [295]. *FilmFinder* generalizes HomeFinder to the task of choosing a movie to watch by transforming the geographical map into a scatter plot with a popularity metric and release year on the axes [2] (Figure 3.5c). An additional type of slider allows the selection of movies from an alphanumeric list of, e.g., directors. *SmartClient* uses a scatter plot to display a subset of options regarding predefined

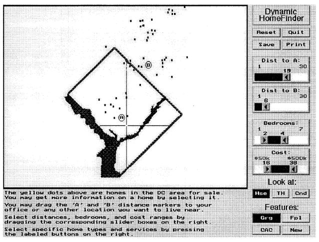

(a) HomeFinder, 1992 [297]

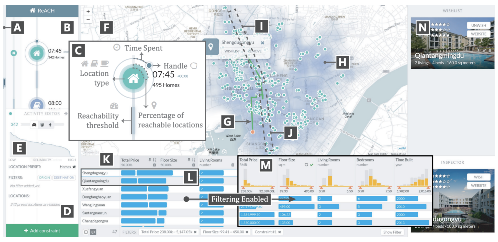

(b) ReACH, 2018 [295]

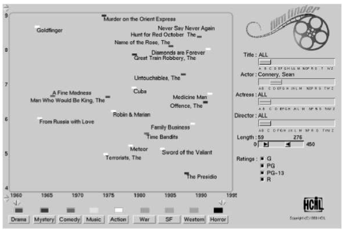

(c) FilmFinder, 1994 [2], as depicted by [156]

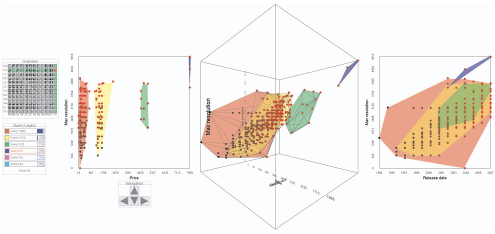

(d) ScatterDice, 2008 [73]

Figure 3.5: Examples for scatterplot-based visualizations to support multi-attribute choices made by non-experts.

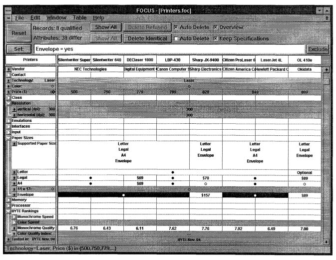

(a) FOCUS, 1996 [257]

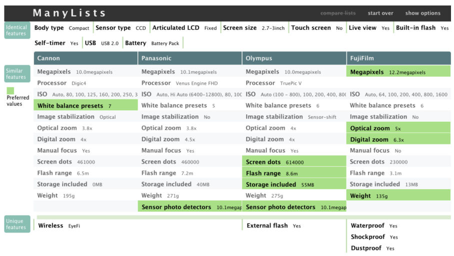

(b) ManyLists, 2012 [149]

Figure 3.6: Tabular visualizations that support generic decisions among products with dynamic querying rather than a weighting approach.

attributes, e.g., flights according to price and total travel time [218]. Constraints on additional attributes can be applied in a tabular or parallel coordinates view. Favorite options can be bookmarked for detailed comparison. For scatter plots to represent more than two original attributes, they have to be arranged in a scatterplot matrix. *ScatterDice* facilitates the navigation of over thousand digital cameras by providing freeform queries in individual scatter plots as well as animated 3D transitions of data and query shapes between neighboring scatter plots [73] (Figure 3.5d). Still, even scatterplot matrices provide a limited perception of complete multi-attribute options.

The house-buying scenario addressed by HomeFinder has also been targeted by an early tabular visualization called *FOCUS* [257]. In another scenario, close to a hundred laser printers and tens of attributes are compressed into a compact table for selection and comparison where products are depicted in columns and attributes in rows. It can be explored by selecting options to be rendered with increased

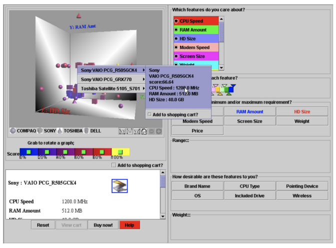

(a) VMAP, 2007 [266]

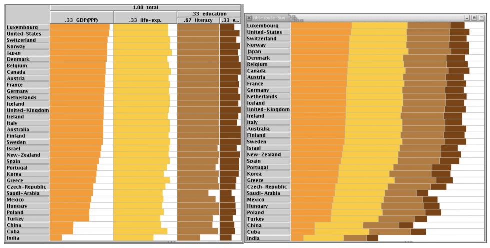

(b) ValueCharts, 2004 [43]

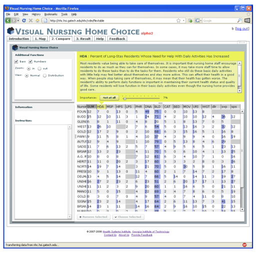

(c) VDM, 2008 [302]

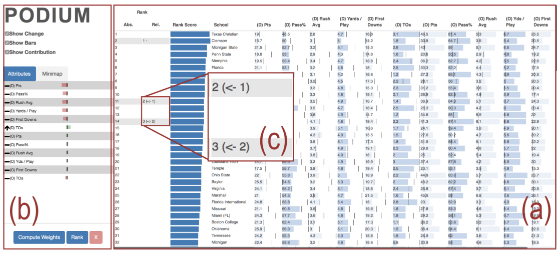

(d) PODIUM, 2018 [284]

Figure 3.7: Most approaches that handle decision-makers' preferences in the form of weight-based scoring are tabular visualizations.

width, expanding and collapsing groups of attributes, and performing queries by clicking on desired attribute values in the table (Figure 3.6a). While FOCUS visualized many options and attributes at once in a compact overview, *ManyLists* aims at an in-depth comparison of up to ten options [149] (Figure 3.6b). Choice scenarios regarding food products and technical products have guided its design. Starting from a table layout similar to FOCUS, attributes are reordered using animated transitions as follows: attributes with identical values across all options are merged in the top row, attributes only available for one option slide down to the bottom of their column, and the remaining attributes are vertically aligned in between.

Besides filter queries, some approaches apply decision-makers' preferences in the form of a weight-based scoring. *VMAP* assigns utility scores to options, laptops in this case, based on the decision-maker's incrementally expressed preferences regarding attribute importance and desirability of attribute values [266] (Figure 3.7a). Many weight-based approaches build upon tabular visualizations (compare Table 3.1). Tabular visualizations are particularly suited to be extended by weight-based rankings because they depict options in an inherent order. A popular example is *ValueCharts* where attribute importance is expressed by resizing the respective column [43] (Figure 3.7b). Collapsed into a stacked bar chart, the columns convey the total utility score of each option as their weighted additive sum. ValueCharts has been applied in the real estate domain or for ranking countries according to their human development index. It has been extended with editors for capturing the desirability of attribute values, which required a rotation of the visualization [24]. Conati et al. empirically

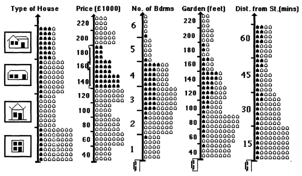

(a) Attribute Explorer, 1994 [280]

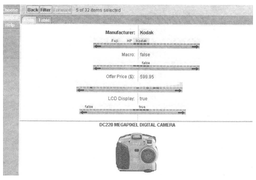

(b) MultiNav, 2000 [141]

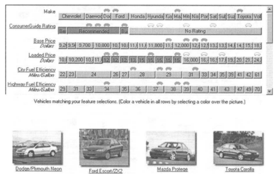

(c) EZChooser, 2001 [298]

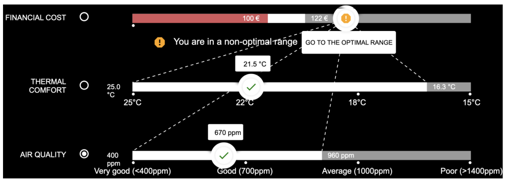

(d) TOP-Slider, 2018 [142]

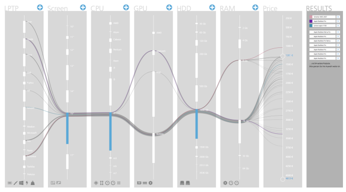

(e) ConfigurationFinder, 2020 [222]

Figure 3.8: Examples for decision support visualizations based on parallel axes, which provide support for interactive querying and filtering.

compared both layouts on a hotel choice scenario and found the layouts to perform differently depending on the decision-makers' visual working memory and visualization literacy [57]. *LineUp* provides more flexible interactions to combine heterogeneous attributes into aggregate scores than ValueCharts [85]. The effects of different weightings can be compared side-by-side using slope graphs connecting the resulting rankings. HomeFinder and ValueCharts also inspired the Decision Map and Decision Table (Figure 3.7c) components of the *VDM* framework for choosing nursing homes [302]. An additional component allows for a weight-based side-by-side comparison of two nursing homes. Weight-based approaches work well for navigating large product catalogs. However, representing preferences by assigning exact weights to each attribute is challenging. To help users better understand the effects of different weights, *Weightlifter* has been introduced [198]. A recent observational user study compared eight interaction widgets for distributing weights to three criteria [116]. Still, the success of weight-based ranking might be threatened by preferences often being "fuzzy, unstable, and inconsistent" [60]. *PODIUM* addresses this challenge by asking users about their preferences regarding items rather than attributes [284] (Figure 3.7d). Still, it remains questionable to what extent decision-makers can be enabled to accurately express their vague and evolving preferences. In view of conflicting attributes, *Imma Sort* provides an alternative to utility-based ranking by sorting options such that multiple attributes exhibit approximate monotonic trends [155]. This helps users predict the values of multiple attributes as they navigate through a sorted list. Its usefulness is demonstrated on scenarios involving the choice of a hotel, movie, and food dish.

Rather than a scoring approach, visualizations that employ parallel axes of different kinds often pursue incremental query construction. Choosing a house to buy has been a common scenario originally addressed by HomeFinder. A more general work took a different approach using linked histograms: the *Attribute Explorer* represents attributes as horizontal histograms that are arranged next to each other [280]. Brushing houses in a certain value range on one histogram axis reveals where these houses appear on the remaining axes (Figure 3.8a). This adds immediate feedback to the to the sliders themselves as opposed to the indirect controls provided by HomeFinder. Perceiving the relationships between attributes has also been the aim of interaction provided in *MultiNav* [141]. It allows users to slide a horizontal attribute axis from left to right, thereby causing the remaining attribute axes to move as well to keep the focused option in the center of the screen (Figure 3.8b). Relationships can be observed from the axes' tendency to move in the same or different directions and at which speed. In *EZChooser*, the horizontal attribute axes are composed of bars representing ranges of attribute values, where the relative count in a bar is reflected in the bar's width [298] (Figure 3.8c). The data set, in this example cars, can be queried by clicking on the value bars in each attribute axis. Among others, the query widgets used in MultiNav and EZChooser have been systematically revisited and informed the design of the *TOP-Slider* [142]. Similar to parallel coordinates without polylines, it uses interlinked parallel attribute sliders to support non-experts in perceiving optimal trade-offs. When a slider handle is selected and moved, the handle is linked to pairs of dashed lines that convey the boundaries of the optimal ranges of the remaining sliders (Figure 3.8d). The *Product Explorer* aims to make standard parallel coordinates accessible to casual users [221]. The authors propose a number of enhancements to increase their usability, including facilitated tracing of product attributes, visualizing gaps in the product data set, and a decision bar that stores exclusive intermediate decisions to free screen space. The most recent interactive visualization for product selection by casual users is a follow-up of the Product Explorer. *ConfigurationFinder* optimizes the parallel coordinates for screen space efficiency by reducing the number of visual elements [222]. It switches from line-based connections to ribbons where applicable (Figure 3.8e), introduces proxy axes to represent groups of semantically related axes and fusion axes to represent combined attribute values of two attributes. The comparison to a popular webshop interface revealed that ConfigurationFinder was perceived as less confusing and easy to use despite its unfamiliar interface. Dimara et al. studied three elementary lossless projection visualizations for their ability to support casual users in making multi-attribute choices: parallel coordinates, a scatterplot matrix, and a tabular visualization [64]. They found no conclusive difference in the decision accuracy between the visualizations. Subjective participant ratings weakly suggest that the tabular visualization, which did not employ any type of utility scoring, is preferred over parallel coordinates. Analyzing the time-on-task revealed some evidence that decisions were made faster with tabular visualizations. A tool that links all three elementary visualization techniques to a map and also supports weighted scoring is *CommonGIS* [12]. It explicitly supports the choice stage (Figure 2.3) of decisions within a geographical context, e.g., choosing counties to receive funding [12] or choosing a skiing resort for vacation [11].

Summary

The commonalities that make multivariate visualization tools particularly accessible to casual decision-makers are 1) the use of well-known visualization techniques like scatter plots, tabular views, parallel coordinates, or maps and 2) the interactive queries and filters, which are often executed via direct manipulation rather than indirect interaction with widgets. Most of the techniques aim at offering a clear overview of the available options, with a few exceptions focusing on in-depth comparisons between a handful of options. Some tools, mostly tabular visualizations, support multi-attribute choice more explicitly by implementing a weighted-sum approach to rank options. While the simple visual encodings help generalize to a range of multi-attribute choice tasks (e.g., different consumer choices), we have observed a decrease of decision tools for the general audience in recent years. We found this observation confirmed by a survey where most tools supporting the choice stage assumed expert users [195]. Our observation might be attributed to the increasing attention received by problem-driven visualization research that often targets expert users. Many of the findings related to tools for casual decision-makers might be helpful for designing visualization tools for professional decision-makers. Still, professional settings might entail different motivations, priorities, environments, or time budgets [104], such that a careful domain characterization and task analysis are needed.

3.2.3 *Professional Decision-Makers: Design Studies*

Professional decision-makers are trained to make decisions as part of their job. Section 3.1.1 has summarized design study methodologies to conduct problem-driven visualization research. While these methodologies can in principle be applied to any user group, design studies are often carried out for domain experts. They have investigated visualization tools meant to support professional decision-making in different contexts. Existing surveys have investigated how and why humans come to conclusions with static visual information [197], how interactive visualization tools assist decision tasks [195], and how such visualizations can be evaluated [64, 71]. While focusing on theory/empirical and design (technique or system) papers, no particular focus has been dedicated to design study aspects. We summarize works on visualization techniques for decision support that present themselves

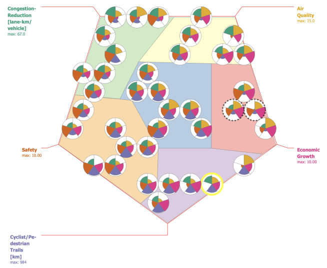

(a) SOMMOS, 2013 [49]

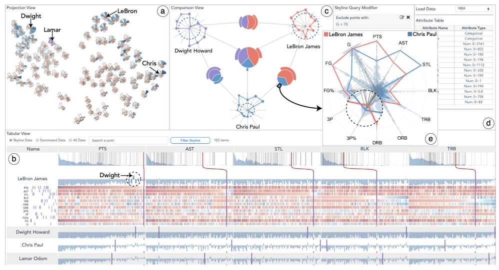

(b) SkyLens, 2017 [309]

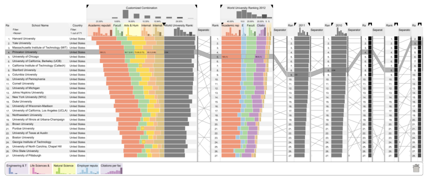

(c) LineUp, 2013 [85]

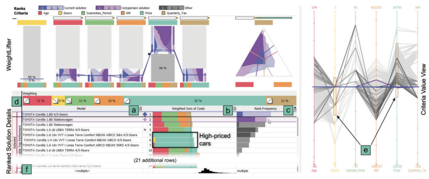

(d) WeightLifter, 2016 [198]

Figure 3.9: Examples of design studies that focus on general-purpose decision-making as a high-level goal.

as a design study or report on a collaboration with domain experts including a detailed data and task analysis or evaluation.

SOMMOS [49], *SkyLens* [309], *LineUp* [85], and *WeightLifter* [198] focused on general-purpose decision-making as a high-level goal. Still, we present them alongside tools for professional decision-makers because they are grounded in literature reviews of scientific domains or collaborations with domain experts. Their extensive capabilities also come with a certain operational and visual complexity, which does not match the requirements of general-purpose tools for casual users. Following Munzner's nested model [185], these works first provide a characterization of decision data, tasks, and requirements, where they focus on commonalities across application domains. On this basis, SOMMOS was designed as an interactive self-organizing map for visualizing a Pareto front, with anchor points corresponding to criteria and radial bar chart glyphs representing individual options [49] (Figure 3.9a). As criteria are considered equally important, the anchor points address the requirement for their objective representation. The glyphs address the requirement for detailed inspection of individual alternatives. Similarly, Zhao et al. use t-SNE for a projection-based overview in their system *SkyLens* [309] (Figure 3.9b). In addition, they dedicate views to the design goals of understanding why an option is Pareto-optimal and what trade-offs are associated with choosing one option over another. Based on a tabular visualization with embedded bars, LineUp allows to create and compare multiple rankings, a requirement they derived from the scenario of investigating changes in university rankings over time [85] (Figure 3.9c). While the parameterization corresponds to a point-wise exploration of the weight space, WeightLifter proposes a global weight space exploration to particularly address the identified mismatch between the need for precise

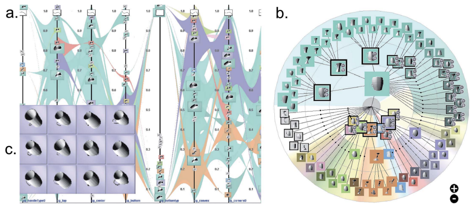

(a) Cupid, 2014 [27]

(b) Dream Lens, 2018 [162]

Figure 3.10: Examples of design studies that aim to aid multi-attribute choice by relating the options' design parameters to their outcomes.

weights and the vague intuition about the relative importance of criteria [198] (Figure 3.9d). The domain-agnostic problem characterizations make SOMMOS, SkyLens, WeightLifter, and LineUp applicable in different applications. This is demonstrated by different usage scenarios, some of them targeting casual decision-makers. Still, additional potential might be unlocked by adapting to domain-specific tasks and requirements. Furthermore, a central projection-based overview like in the case of SOMMOS and SkyLens requires additional representations to ensure full visibility for multi-attribute choices that require (all) attribute values to be directly retrieved [64]. In addition, ranking approaches like LineUp and WeightLifter, which combine attribute values into weighted scores, require significant cognitive effort from decision-makers, who need to capture their vague and at times interacting preferences by single precise weights.

Vismon was introduced as a trade-off analysis tool for policy-making [35]. In contrast to the works before, the design of Vismon built upon a characterization of domain-specific data and tasks, namely in fisheries management. At its core, contour plots are used to depict the value distribution for each criterion as a function of the two design parameters. A few selected options can then be compared in a trade-off view, where bar charts by default show each criterion across the options in focus. Still, Vismon does not provide an overview of alternatives with all their properties. Tools to visualize and explore alternatives in the context of generative design allow for considering qualitative aspects such as geometry along with quantitative criteria. *Cupid* makes use of superimposition to allow users to relate the abstract parameters of a geometry generator to the resulting three-dimensional shapes [27] (Figure 3.10a). In contrast, *Dream Lens* juxtaposes views for the shapes and parameters of the designs [162] (Figure 3.10b). The above design studies approach decision-making by navigating through given design options and asking about their outcome. Multi-attribute choices often require the reverse: given a desired outcome, what does the design option need to look like?

As there exists no direct inverse mapping from outcome to design parameters, an interactive exploration of varying design parameter settings and their outcomes is needed (Section 3.2.1). Matković et

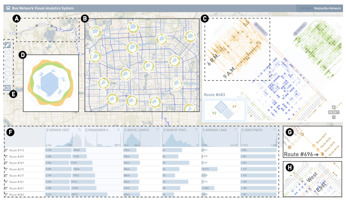

(a) Weng et al., 2020 [294]

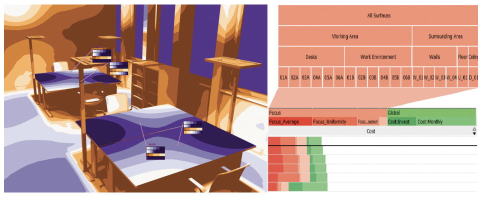

(b) LiteVis, 2015 [254]

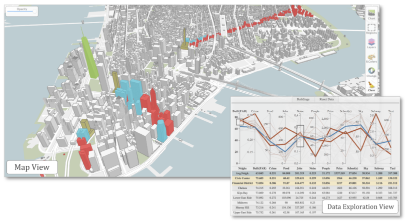

(c) Urbane, 2015 [76]

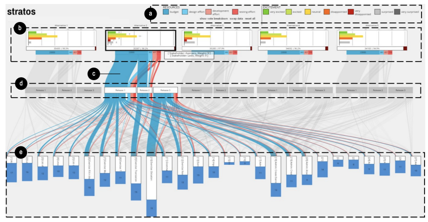

(d) STRATOS, 2015 [14]

Figure 3.11: Examples of visualization support for decision problems that involve information about temporal and/or spatial aspects.

al. combine scatter plot, histogram, and parallel coordinates to assist engineers in analyzing and understanding the parameter space of fuel injection systems [163] (Figure 3.3a). Their basic information visualization techniques have proven valuable for gaining a broad understanding of the system behavior. However, the case study only covers the identification of desirable regions in the parameter space and does not ultimately proceed until a final choice remains. Grounded in the same use case, Berger et al. extend the inspection of a limited number of sample points to a continuous navigation of the design space while observing the behavior of multiple decision criteria [28]. Experts particularly appreciated the prediction of combined criteria changes in a local neighborhood when increasing or decreasing design parameter values (Figure 3.3b), which significantly helps the identification of potential for improvement or, where necessary, for trade-offs. However, their approach centers around additional sampling. It thus falls into the category of interactive approaches, while we consider multi-attribute choice as an a posteriori decision problem where all options are known in advance (Section 2.2.1).

Besides strictly multivariate options, design studies have also investigated decision support in spatial scenarios. *RelEx* resulted from a design study on optimizing the traffic flow in communication networks [239]. The visual network analysis tool supports engineers in modifying an existing communication network such that physical constraints are met and costs for signal routing are optimized. Similarly, Weng et al. study the generation and evaluation of bus route candidates in a complex bus network to meet constantly changing travel demands [294] (Figure 3.11a). With *LiteVis*, Sorger et al. propose a visualization that helps simulate and evaluate different lighting setups for buildings or open spaces [254] (Figure 3.11b). Similarly based on a three-dimensional scene, *ManyPlans* allows for the generation and evaluation of response plans with protective measures for different flooding scenarios [291]. In both cases, decision criteria can be both multivariate, e.g., investment cost, and spatial, e.g., illumination or water level at different locations in a scene. Like LineUp [84] and WeightLifter [198], the approaches use a weighting approach to compute a ranking of the generated solutions, which is visualized as a stacked bar chart. While ranking computation is a valid approach to handling options in multi-attribute choices, it relies on the challenging task of establishing a meaningful weighting of decision criteria in advance. *Urbane* also targets a three-dimensional scene, but exposes urban development options in a parallel coordinates plot instead of a ranking approach [76] (Figure 3.11c).

Besides spatial data, decision-making might also involve an analysis of temporal or spatio-temporal data. *STRATOS* allows for deciding among a handful of software release plans, depicting stakeholder satisfaction and feature coverage alongside the flow of resources through a hierarchy of software releases and features [14] (Figure 3.11d). In clinical environments, visual assistance tools support decisions for diagnosis or treatment planning, e.g., based on event sequences [148, 184] or blood flow data [204]. To help choose among possible marketing activities, Guo et al. visualize event sequence predictions of likely customer behaviors following a certain action like sending an email [87]. Due to the complexity of individual options introduced by the temporal and spatio-temporal dimensions, these approaches only allow for a comparison among a handful of options.

3.2.4 *Interaction*

Like decision-making can be a goal of visualizing data, it can also be a goal of interacting with a visualization. Interaction is essential for visual exploration tasks and likely also for decision tasks [64]. It occurs at all stages of the visualization pipeline (i.e., data transformation, visual mapping, and view transformation) [42], where it serves different data-oriented intents [66]. On a high level, interactive manipulation allows decision-makers to acquire multiple perspectives on the data and form insight. On a low-level, it allows decision-makers to control (the representation of) the solution space. Interaction generally encourages decision-makers to engage into thinking and reflecting on the data being explored.

We briefly summarize interaction techniques according to their key role in a visualization-based decision process. We take a task-centric approach where we loosely group the interaction techniques by *why* they are used and describe within each group *how* an interaction mechanism helps the respective task. Note that this is neither an exhaustive collection of interaction mechanisms nor an exclusive assignment of mechanisms to tasks.

Being confronted with a (large) number of alternatives and attributes, decision-makers benefit from an initial overview of the entire data set to get a feeling for the extent of the design space and the range of achievable performances. Interacting with such an overview marks the beginning of an explore task where decision-makers search for alternatives with particular (desired) characteristics [38]. This might involve changing how the information is represented [66], such as adjusting the color-coding of alternatives to observe the distribution of decision criteria values or derived attributes like a computed utility score [38]. Mouseover highlighting is useful to gain a first impression of the characteristics of alternatives or of the frequency and distribution of attribute values. It can also reveal details of aggregations [38] that might have been introduced to roughly depict the characteristics of the solution space using a manageable number of representatives. Interaction at this stage might also involve navigating the information [66], for example via zooming, panning, or scrolling, which manipulate the visible viewport but leave the visualized data unchanged.

Multi-attribute choices center around conflicting attributes that need to be balanced. Recognizing relationships between attributes or alternatives that hint at required trade-offs is thus essential. Many interaction mechanisms support this task by changing the layout of a visualization [66], e.g., by re-arranging the representations of attributes or alternatives. For example, attributes might be re-ordered to perform side-by-side comparisons, or sorted to express their relevance to the choice task [64]. Some visual encodings, e.g., tabular views, also allow for re-ordering of alternatives. Sorting according to ascending/descending attribute values, derived utility scores, or similarities is a simplified variant of alternative reordering [64]. Scanning through the values of a sorted attribute via mouseover and simultaneously observing the highlighted values of another attribute might also help in discovering attribute relationships. Similarly, brushes can be created (and moved) to highlight a subset of alternatives with respect to one or more attributes and observe their distribution on the remaining attributes [26]. This is also one application of the focus-and-context technique, where those parts of the data, which are currently in focus, are visually discriminated from all the rest (the context) [92]. When used in conjunction with linking, brushing elements in one view consistently highlights the same (brushed) alternatives in all other views, thus allowing users to understand relationships across different perspectives [225]. Where attributes in fact exhibit relationships that introduce redundancy to the choice task (e.g., a positive correlation where two attributes can be optimized simultaneously), one of them might be hidden.

Section 2.1.4 showed that decision strategies, in particular the elimination by aspects strategy, involve filtering alternatives into acceptable and unacceptable regions to narrow the solution space. Again, this is an explore task in the sense of searching for targets with matching particular (acceptable) characteristics [38]. Filtering is performed by creating brushes that represent desired attribute values, where undesired alternatives not satisfying the constraints are filtered out and hidden from the display [26]. As a choice typically needs to satisfy multiple constraints and preferences, several brushes can be combined to a composite brush using the logical AND operation [288]. This puts alternatives relevant to the choice into focus and simplifies their distinction from those that are less desired. Unlike multivariate analysis problems, decision problems usually involve additional information about whether certain attributes should be minimized or maximized. This knowledge could be incorporated into the design of brushing mechanisms to reduce the set of operations needed.

Once decision-makers have reduced their choice to a manageable number of multi-attribute options, making a final decision requires a careful consideration of their detailed characteristics. This corresponds to the query tasks of identifying the characteristics of one alternative and comparing it to those of another alternative [38]. For this purpose, mouseover highlighting of an alternative can be extended by tooltips depicting the raw attribute values of the hovered alternative as details-on-demand [64]. However, these information disappear once the mouse leaves the alternative. In contrast, bookmarked options are persistently visible, even if they are not part of any other selection [38]. This enables a direct comparison with respect to each of their attribute values. To help the in-depth comparison among multiple favorites, alternatives might also be annotated with labels storing, e.g., their most remarkable characteristics. Similarly, to keep track of the decision process, intermediate decisions might be annotated with, e.g., the reasons for exclusion of an alternative.

3.3 RESEARCH GAPS

In Section 1.2, we identified research challenges for assisting data-informed decision-making with visualizations. Here, we conclude with a summary of research gaps that remain from the challenges after reviewing the state of the art in Sections 3.1 and 3.2 (Figure 3.1). In the course of this thesis, we will transform these gaps into research goals, describe our concept to address them, and demonstrate its applicability in different use cases related to multi-attribute choice.

Prior works on how visualizations can support decision-making are highly valuable. Still, despite decision-making being researched as a central goal of visualization, important gaps remain to be filled, which we list below. We refer to the research challenges that each gap emerges from (i.e., which are not yet solved) in brackets. A main observation from our literature review is that there are few works that:

- **Systematically elicit tacit knowledge and strategies (C_{CHA}) –** Although an effective decision support requires an integration of explicit data with implicit knowledge [226], few works report on the procedures they followed to extract the expertise and experience involved with their users' problem-solving. No prescriptive steps are available yet to guide visualization researchers in eliciting the tacit knowledge involved with decision-making. Instead, most domain characterization reports focus on the outcome, i.e., data and task abstractions. While ethnographic methodologies provide systematic approaches to knowledge elicitation, the difficulty is to transfer the practices from their limited scope to contexts that involve large data and technical artifacts.
- **Explicitly distinguish decisions from analytic tasks (C_{CHA}) –** Although decision tasks are different from traditional analytic tasks [195], most works that are meant to support decision-making actually center around analytic tasks. Little experience has been gathered with developing a meaningful outline of a decision problem. Raising awareness for what can or cannot be considered a decision task could help conduct more rigorous research of visualization-based decision support. The difficulty is to establish a clear terminology and abstraction concept around decision tasks in order to disambiguate research claims and move the focus of visualization research from analysis tasks to decision-making activities.
- **Tailor visualizations to constructive preferences (C_{SUB}, C_{VIS}) –** Although multi-attribute choices are ill-defined problems and do not have an optimal solution, few visualization designs explicitly target tasks like constructing and validating preferences, comparing the gains and losses of options, and reconciling conflicting information to make a final choice. Multivariate visualizations serve as a solid foundation for basic analytic activities but fall short of determining an alternative's value to a decision. The simplified functionality of tools that target casual decision-makers might not meet the flexibility needs of experts. Tools designed to support expert decisions miss different core aspects of the targeted choice problem. In particular, approaches proposing (weighted) scores to establish an order among alternatives cannot fully represent the preferences and requirements of decision-makers, e.g., by single weights. The difficulty is to fuse the tacit knowledge and preferences of decision-makers with the explicit data to provide an environment, in which decision-makers can discover what is important to their choice.
- **Target decisions where trade-offs affect each other (C_{VIS}) –** Although most decisions taken will influence their surround-

ings, dependencies between decisions have not been addressed so far. To date, no work exists that helps decision-makers perceive multiple alternative sets associated with such problems. Visualization approaches have already linked data sets stemming from different acquisition modalities or simulation models [122]. However, such data compositions have not yet been viewed in the context of decision-making. The difficulty is to generate a coherent representation of such data, to navigate the large number of possible outcome combinations, and to judge the overall goodness as an emergent property that the individual choices do not exhibit when viewed separately.

- **Are validated on expert choices in real-world settings (C_{EVL})** – Although many works mention decision-making as a main goal, they rarely report on case studies that assess whether domain experts have been indeed helped in making their choices. Several tools are validated with hypothetical experts or on simplified decision tasks. Increased speed is also often considered as a quality indicator for task performance. But in decision-making, it is more important that decision-makers obtain a differentiated, and therefore trustworthy, understanding of their options rather than to quickly jump to conclusions. The difficulty is that improved decision quality is more difficult to grasp due to a lack of ground truth (compared to low-level analytic tasks) and that domain experts who were not involved in the design process might be rare. Furthermore, conducting and analyzing observational studies demand a significant amount of resources, both from domain experts as well as visualization researchers.
- **Observe long-term benefits based on self-initiated use (C_{EVL})** – Visualization tools are typically evaluated with respect to their short-term usefulness only. Few works (whether or not targeted at decision-making) follow-up on their tools' long-term usage, let alone adoption, in the target domain after a certain time period. Besides the general challenges associated with case studies, the difficulties of studying adoption can mostly be attributed to logistics: the tool needs to be promoted among target users, self-initiated use needs to be noticed by the researchers, and checking back after some time might rely on a continued collaboration.
- **Focus on design studies for multi-attribute choice (C_{HCD})** – Although decisions are highly subjective, few works ground the design of their visualization tools in close collaborations with professional decision-makers. Many visualization papers focus on decision tasks of narrow complexity that involve small data sets [67], often studied in (controlled) artificial settings. When designing for real-world settings, the difficulty is to perform an accurate domain characterization and abstraction against the background of an ill-defined decision problem.

Part III

MAIN CONTRIBUTIONS

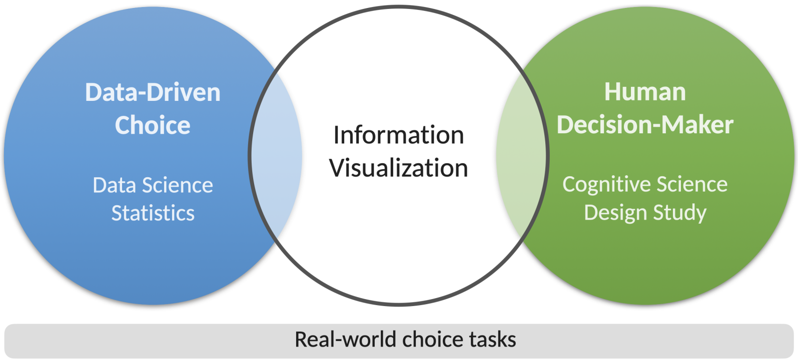

Figure 4.1: Grounded in real-world applications, research on making choices with visualizations touches both data science and human science.

4 CONCEPT FOR VISUAL ANALYSIS SUPPORT FOR MULTI-ATTRIBUTE CHOICE

THE central research question targeted in this dissertation is how we can design and validate interactive visualizations to effectively assist experts in making real-world choices among many multi-attribute alternatives. To this end, we need to bridge the gap between decision theory focusing on controlled settings of narrow complexity and visualization research focusing on real-world exploratory settings by analysts instead of decision-makers. In this chapter, we describe the conceptual approach taken in this thesis, i.e., the relevant ideas and rationales that guided our research.

What the reader can expect from this chapter:

- A contextualization of multi-attribute choice in the scope of data science, human science, and visualization (Sec. 4.1).
- An outline of the research goals we address (Sec. 4.2).
- A characterization scheme to abstract a given choice task using pairs of data, user, and task properties to help visualization researchers precise their decision-support claims (Sec. 4.3).
- An instantiation of the characterization scheme that results in an operational definition of our targeted choice task (Sec. 4.3).
- A qualitative research design for the visual analysis for multi-attribute choice. The concept considers the researchers' visualization design process and the users' decision process as orthogonal dimensions (Sec. 4.4).

4.1 RESEARCH SCOPE

Assisting multi-attribute choices with interactive visualizations is multi-faceted and thus requires a multidisciplinary approach. We provide a contextualization of the topic by relating it to three research areas (Figure 4.1): (1) the research problem is about making *data-driven choices*, (2) the data analysis requires judgments of a *human decision-maker*, and (3) the presented solutions make use of *information visualization*. Real-world applications from different target domains serve as information sources.

When making a choice between options, we are dealing with multivariate data items that can be organized in a data table [193]. In this data table, the rows represent options in the form of data items and the columns represent attributes [64]. It thus makes sense to consider known ways to extract patterns and meaning from such data. Such approaches are researched in the field of data science (Figure 4.1, left).

At the same time, deciding between options involves human judgment. This includes, among others, the role of prior knowledge, cognition, (rational) reasoning, and decision strategies. In an attempt to understand and build upon how humans make choices, we can borrow from the domain of human science, in particular from the fields of decision theory and cognitive science (Figure 4.1, right).

Finally, understanding the information that is hidden in the data to make subjective judgments relies on visual representations that facilitate interpretation. We propose information visualization approaches that build upon both disciplines to provide human decision-makers with an effective means to make informed choices (Figure 4.1, center). These choices arise from their personal or professional lives. The visualization research presented in this thesis highly values the work with real users to solve their real-world problems. In this way, the needs, tasks, and goals framing the decision-making in a particular domain can be carefully considered (Figure 4.1, bottom).

4.2 RESEARCH GOALS

Surveying existing works against the background of the research challenges outlined in the introduction (Section 1.2) revealed a number of research gaps (Section 3.3), which we will address in the course of this thesis. From the fundamental research challenges and identified gaps, we derive a number of research goals (Figure 4.2) to motivate the conceptual framework and contributions:

- **Characterize multi-attribute choice** – The related works presented in Section 2.1.1 showed that the task of making a decision is described differently among research disciplines. The wide spread of decision-related tasks is a strong motivation for our research on multi-attribute choice. We identified a number of

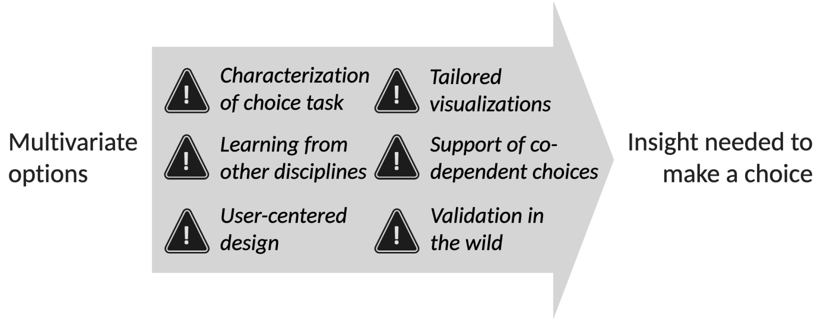

Figure 4.2: How can interactive visualization assist and inform choices? Aiming to connect multivariate data to insights that are relevant for a choice raises a number of research goals for analysis support.

commonalities that can be considered a basic cross-disciplinary understanding of decision tasks. Still, they remain high-level and mainly refer to the information base (i.e., alternatives, attributes, and goodness indicators). Further environmental influences relevant to the design and evaluation of visualization support are not covered. A characterization of multi-attribute choice to inform visualization design in this work should include not only the data aspect but also the perspectives of users and tasks [176].

- **Learn from other disciplines how to study decision tasks** – The background on decision theory (Section 2.1) showed that the field has devoted a long history to studying decision-making and its cognitive foundations. Its findings are likely useful for studying decision tasks in visualization research. Yet, they are rarely considered [67]. Models and concepts from such disciplines could serve as a starting point in moving towards general abstractions of decision tasks and decision-makers, upon which visualization design and evaluation can build. Domain characterization, requirements specification, and task analyses should thus be viewed from a multidisciplinary perspective.
- **Apply and report on user-centered design** – Visualization designers may choose from a variety of visualizations and designs. Some create incorrect conclusions, most do not really help, and only some of them are useful for the task at hand [186]. We thus need to choose visualizations depending on the specific purpose they are meant to serve. With the subjective contexts of decision tasks in particular, approaches to visualization design need to be user-centered. Similarly, Section 2.1 showed that decision theory has moved from rational models and controlled experiments to empirical models and observational methodology. Still, Section 3.2.3 showed that research from a design study perspective providing a contribution towards solving a real-world decision problem is rare. Moreover, professional decision makers have found to be under-supported by visualization research [195].

Future efforts in designing visualization support would benefit from targeting professional decision-makers and from being grounded in a close collaboration with domain experts.

- **Consider constructive preferences in visualization design** – Designing effective visualizations for choice tasks requires a trade-off between information volume, simplicity and visibility [71]. This trade-off needs to be made against the background of decision-makers constructing their preferences on the spot, rather than drawing from well-defined pre-existing preferences. In doing so, they utilize a variety of approaches that they might adapt depending on the varying context and task demands during the course of making a decision. Visual representations and interaction mechanisms need to be carefully designed, not only to be well understood, but also to provide a flexible environment for the incremental construction of subjective judgments.
- **Support co-dependent choices** – Many decisions are made among options that refer to single units like an apartment to rent or a mechatronic component to take to production. However, other decisions involve multiple choices that affect each other, be it mechatronic components that are operated together [161] or life partners searching for jobs within a reasonable commuting distance from each other. The core challenge of making trade-offs extends beyond one single choice. Visual representations need to make these side effects visible to support decision-makers in subordinating individual choice goals in favor of those of the overall decision. Given the general significance of such decision problems, visualization research would benefit from broadening the profile of targeted tasks towards such co-dependent choices.
- **Validate visualization designs in the wild** – Any visualization design needs to be validated with evidence that it does indeed help solve the targeted decision problem. This requires study conditions that are closest to the real-world practice of a domain expert. It can be achieved with field studies, where free tool use by domain experts in a real-world setting is observed. Lab studies with domain experts tend to provide less rich findings [243]. In line with the call for more user studies that contain decision tasks [67], Section 3.2.3 found that few works report on their tool's usefulness in case studies with domain experts. Even fewer works report on self-initiated long-term use or adoption. Problem-driven visualization research would benefit from an assessment of decision support tools in the wild, both in the short-term and in the long-term.

4.3 CHARACTERIZATION OF MULTI-ATTRIBUTE CHOICE

Before we can frame a concept to approach the research goals, we need to precise our understanding of the targeted decision problem. Table 2.1 revealed significant commonalities regarding the decision activity, alternatives, attributes, and goodness indicators of decision tasks across disciplines. Most tasks deal with a variation of choosing among a set of multi-attribute alternatives, aiming to maximize some value. The most recent definition in the field of information visualization refers to multi-attribute choice as the task of "*finding the best alternative among a finite set of alternatives, where alternatives are defined across several attributes*" [64]. We opt for "most preferred" over "best" to indicate that some subjective or preferential value that is not necessarily consistent with rationality is to be maximized.

Section 1.2 illustrated that multi-attribute choice is an ill-defined problem involving manifold alternatives, conflicting demands, shifting goals, implicit knowledge, and potentially high stakes (Figure 4.3). The same section identified the need to take a closer look at real-world decisions made by real users. In this section, we detail the properties of multi-attribute choice to clarify the different guises in which such a task might come. The ultimate goal is a clearer understanding of the task as it is targeted in this thesis.

4.3.1 *Properties of Multi-Attribute Choice*

By reading through literature and talking to domain experts over multiple years, we gained an overview of the wide range of situations and disciplines, in which choices are made. Based on this experience, we categorize the properties of the examples we encountered to develop a systematic view on the diversity of real-world choices. The goal is to identify characteristics that help describe the multi-attribute choice task to design for more explicitly.

Our characterization is grounded in three practical questions:

- What **data** inform a decision?
- Who are the **users** making a decision?
- What conditions frame the **decision task**?

These questions correspond to data, users, and task as the three main aspects to consider when designing visualizations [176]. In the following, we will provide detailed explanations of each aspect, including characteristics, examples, and their implications for visualization design. Multi-attribute choice preserves the following properties:

Data

- **Number of options: few vs. many** – The fewest options would be a binary decision, e.g., whether to go for a walk or not. A

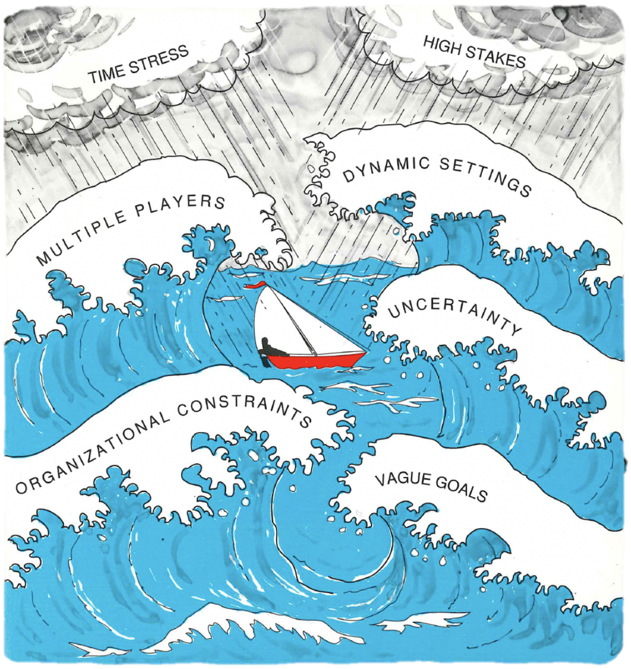

Figure 4.3: Choices in natural settings are difficult. In this work, high stakes, dynamic settings, uncertainty, and vague goals are particularly relevant. Image reprinted with permission from Klein et al. [130]. Permission for reuse must be obtained from the copyright owner.

typical example for the availability of many options are consumer choices, e.g., which of the hundreds of hotels in a city to book for a vacation. Visualizing many options can require advanced aggregation or focus-and-context techniques to address visual clutter and support user navigation through the solution space.

- **Number of attributes: few vs. multiple** – A decision involving few yet conflicting criteria might be an investor balancing risk and return indicators of stocks to buy. In an architectural context, the variety of attributes might range from geometric and material parameters over mechanical stresses to energy efficiency. The number of attributes that can be effectively displayed depends on the chosen visual encoding.
- **Availability: known vs. progressive** – Decisions are often made among a fixed set of solutions, e.g., from a product catalog or pre-computed simulations. For other decisions, new solutions in a region of interest can be generated progressively during the decision process, e.g., an operator steering the optimization of wastewater treatment. With a progressive steering, the focus shifts from multivariate visualization of options to the interaction between the decision-maker and an optimization algorithm.
- **Alternatives: comparable vs. non-comparable** – Many decisions are among comparable alternatives that share the same attribute space. Still, one can also face a choice, in which the attributes

defining the options differ, such as deciding whether to buy a new camera body or a camera lens. Non-comparable data points with partially disjoint attribute sets significantly complicate the visual mapping.

- **Consequences: certain vs. uncertain** – Choices are often made based on expected consequences. As an example, a decision-maker might be certain about the horsepower of a car, while other properties might be subject to uncertainty, e.g., its reliability if the car has been launched only recently. Visualizations meant to serve as a decision-making tool need to communicate uncertainty information to avoid bias.

Users

- **Decision-maker: casual vs. professional** – Casual decision-makers face smaller or larger decisions at some points in their personal lives. Professional decision-makers are experts, whose job it is to repeatedly make decisions within their field and who are trained to do so. Understanding the motivation, priorities, environments, and cognitive processes of the decision-maker helps derive appropriate requirements for visualization design.
- **Authority: oneself vs. others** – Personal decisions are usually made on behalf of oneself and do not require justification beyond that. Expert decisions are typically made on behalf of other stakeholders, e.g., mechatronic scientists who design electric drives for their customers. As expert decision-makers are answerable for their decision, visualizations can help them articulate the rationales of their choice and verify the choice together with the additional stakeholder.
- **Social setting: individual vs. group** – Individual choices involve only the knowledge, intuition, and judgment of the decision-maker. Collaborative decisions are often made in organizations and require consensus among stakeholders, e.g., politicians from different parties deciding about a policy to be implemented, researchers choosing a best paper award recipient, or multidisciplinary tumor boards planning patient treatment. In a group setting, all relevant stakeholders should have been heard. Visually comparing the individual stakeholder's preferences can support consensus building.
- **Involvement: high vs. low** – Decisions that are based on ill-defined, incomplete, competing, or evolving goals require greater involvement and thought by a decision-maker to exploit the available room for trade-offs. This motivates the use of visualization for a human-in-the-loop approach that enables decision-makers to apply their knowledge and experience. Where goals are clear and stable from the beginning, decisions can be based on rules or on automation, e.g., assigning teachers to school lessons.

Decision Task

- **Stakes: low vs. high** – Stakes refer to the severity of a decision's consequences. An example for a choice where the stakes are low is which book to borrow from a library, where a bad decision can be easily reverted. High stakes are associated with many decisions in medical treatment planning, e.g., which patient should receive an organ. In the most critical case, the health or even survival of a human depends on such decisions. Low-stakes decisions in a casual setting require usable rather than sophisticated visualizations. High-stakes decisions likely require a high level of data comprehension and application of domain expertise, posing significant demands on the analytic features and visual encodings of a decision support tool.
- **Time frame: critical vs. not critical** – Some decisions need to be made under time pressure, such as a fireground commander being called to an emergency and having to decide whether to initiate search and rescue. In other cases, the available time frame matches or even exceeds the time required to make a decision, e.g., an urban planner tasked with the reduction of traffic bottlenecks. Accordingly, the time and incentive decision-makers bring to the table to take advantage of visual encodings might vary as does the required level of amplified cognition.
- **Frequency: routine vs. ad hoc** – Routine decisions are choices that we make repeatedly under similar conditions, e.g., whether to have coffee or tea in the morning. In contrast, many life choices are non-routine, ad hoc decisions that pose individual cases, such as buying a house or choosing a university to attend. The decision frequency likely influences the amount of time decision-makers devote to a data analysis. Again, few time resources require a visualization support to be easy to learn while one might assume a familiarity for routine decisions repeatedly made with the same decision support tool.
- **Procedure: prescriptive vs. spontaneous** – Some clinical decisions prescribe particular steps, e.g., a decision tree that queries the (non-)occurrence of symptoms. Spontaneous decisions do not involve imposed instructions but might rely on characteristic steps that are repeated as necessary. Decisions without any imposed strategy benefit from visualization support for an exploratory analysis that helps decision-makers understand on which information they base their decision.

4.3.2 *Targeted Choice Task*

We use the above characterization scheme to define the multi-attribute choice task that is targeted in this thesis (Table 4.1).

Table 4.1: Overview of the characterization scheme for multi-attribute choice tasks. The highlighted properties mark the characteristics of the choice task as targeted in this thesis.

Data	Number of options	few vs. many
	Number of attributes	few vs. multiple
	Availability	known vs. progressive
	Alternatives	comparable vs. non-comparable
	Consequences	certain vs. uncertain
Users	Decision-maker	casual vs. professional
	Authority	oneself vs. others
	Social setting	individual vs. group
	Involvement	high vs. low
Decision Task	Stakes	low vs. high
	Time frame	critical vs. not critical
	Frequency	routine vs. ad hoc
	Procedure	prescriptive vs. spontaneous

The data underlying our targeted choice task involves *many*, i.e., multiple tens to hundreds, alternatives. Each alternative is defined across a set of *multiple* attributes. We expect some of these attributes to constitute conflicting decision criteria, but which attributes this actually applies to is not necessarily known beforehand. The attribute set might or might not include additional design parameters describing the configuration of an alternative. We assume the set of alternatives to be fixed and *known* in advance, without caring about how it comes to be defined or how it might change with experience [154]. We use the term "fixed" over "finite" to indicate that the set of alternatives is not only bounded in size but also immutable. The immutability extends across all alternatives, meaning that the attribute set is the same, i.e., *comparable*, for all alternatives (although we will also come across a new type of choices that will involve *non-comparable* alternatives). We expect all attribute values for all alternatives to be available. We further assume that the consequences represented by attribute values can be considered *certain*.

Our research targets *professional* decision-makers who are interested in applying their domain knowledge, experience, and expertise in choosing the most preferred course of action. In most application domains, professional decisions are made on behalf of *others*, meaning decision-makers are answerable for their decision and must be able to articulate the rationale of their choice. We address *individual* decision-making rather than making decisions collaboratively in groups. Expert decisions are likely to come with a *high* involvement where vague

Figure 4.4: A core part of our qualitative research approach is the close collaboration with domain experts.

goals and dynamic environments are met with thoughtful trade-offs that exploit the room for criteria improvement.

Another consequence of our focus on expert decisions is the tendency of decision tasks towards *high* or medium stakes that require careful compromises. We consider decision tasks that are *not time-critical*, i.e., where decision-makers can and do devote enough time to making a decision. In line with this, we focus on supporting *ad hoc* decisions on a case-by-case basis. Accordingly, we research how visualizations can assist *spontaneous* decision-making that might involve but does not prescribe a particular procedure.

4.4 RESEARCH DESIGN

Given our research goals (Section 4.2) and targeted decision task (Section 4.3.2), this section details the strategy for studying the assistance of expert choices by interactive visualizations. Our research approach will involve an analysis of real-world choices made by experts, the design and validation of visualization artifacts to help them select the most preferred option, and reflections on how our lessons learned might affect visualization research.

Decisions are inherently subjective. We assume that preferences and decision rules are context-dependent and therefore developed at the very moment they are needed during the course of making a choice. In other words, decision-makers learn what they want or need as they explore the available options. This conception means to consider multi-attribute choice as a constructive problem (compare Section 2.1.5). To design and validate effective visualizations that assist decision-making, we need to understand how humans in a particular real-world setting make choices with (and without) visualizations. Aiming to understand people and their experiences in open-ended real-world settings suggests the use of a qualitative approach. This also shows in decision theory, where approaches have moved from rational models and controlled experiments to empirical models and observational methodology (Section 2.1). In line with other disciplines studying human behavior, we therefore engage in *qualitative field work*. This approach offers insights into people's perspectives on decision-making practice on the basis of data gathered through observations and interviews (Figure 4.4).

A wide-spread methodology involving qualitative field work in applied visualization research are design studies. Their focus on designing visualizations for real-world problems and on engaging with domain experts (Section 3.1.1) aligns closely with our research goals (Section 4.2). We therefore approach the topic of visual analysis for multi-attribute choice from a *design study perspective*. Some researchers might question the scientific validity of this perspective for not being reliable, replicable, or generalizable due to missing objectivity. In contrast, we refer to the argumentation by Meyer and Dykes about scientific rigor in visualization design studies [168] and suggest to embrace subjectivity as a strength for constructing new knowledge through interpretation of qualitative data.

Information gathering in a design study can tap different sources: 1) reading domain literature, 2) asking domain experts about tasks, practices, and challenges, and 3) observing them on real-world decisions in a think-aloud manner. Abstractions of the decision problem and data are needed to assemble individual design study results into a big picture. The technical contributions of this thesis will be interactive visualizations to be used by expert decision-makers to choose among large numbers of multi-attribute alternatives. While novelty is considered important, the focus will be on leveraging existing visualization and interaction techniques to provide an effective solution for the given real-world decision. Validation with real decision-makers is needed to assess how well a visualization assists their professional decisions. It should help understand how visualization supports the generation of actionable insight and how decision-makers interact with it in a real-world setting. Following commonly used methods in qualitative research, our validation will build upon observational methods as well as interviews and case studies.

Generally, design studies that target real-world decision problems involve two perspectives (Figure 4.5): the perspective of the decision-maker who works through a decision process and that of the visualization researcher who works through the visualization design process.

The *decision process* refers to the steps that decision-makers follow when approaching a decision problem (Figure 4.5a). From decision theory, we know how humans make decisions. Section 2.1.3 introduced an influential three-stage decision process model by Herbert Simon [252] (Figure 4.5a, bottom). Given our conception of multi-attribute choice as a constructive problem, we refine the decision process axis by aligning Simon's model with a process highlighting the role of

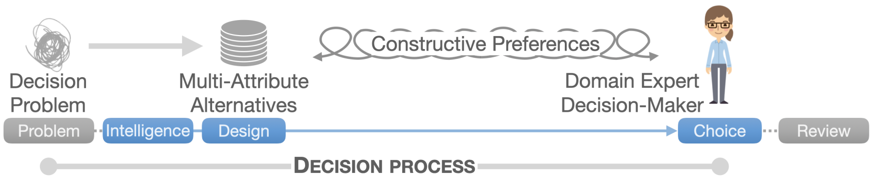

(a) Decision Process

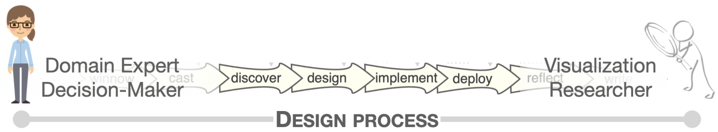

(b) Design Process

Figure 4.5: Two perspectives are relevant for our concept: (a) the decision process of a decision-maker, for which we align the three-stage model [252] with the idea of constructive preferences [32] and (b) the design process of a visualization researcher, which we instantiate with the core phase of the nine-stage framework [243].

implicit constructive preferences [32] that are applied to the explicit alternatives set (Figure 4.5a, top).

The *design process* refers to the steps that visualization researchers follow when designing and validating visualizations that help solve real-world problems (Figure 4.5b). From research on design studies, we know how to approach problem-driven visualization design. Section 3.1.1 presented different models that provide guidance on how to systematically conduct a design study. For our purpose, we need a model that is grounded in experience with real users and their problems and that offers practical advice on this basis. As a representative for the design process, we therefore consider the nine-stage framework that originated from recurring steps across many design studies [243]. More precisely, we focus on its four-stage core phase that covers the actual execution of a design study.

Section 4.2 emphasized the relevance of decision tasks as a subject to be studied by problem-driven visualization design. For visual decision support to be effective, the decision-maker's decision process needs to be considered by the visualization researcher's design process. Our concept targeting visual analysis for multi-attribute choice thus combines both processes, design process and decision process, as orthogonal dimensions. Figure 4.6 shows the concept making use of two display axes. The axes span a two-dimensional conceptual space, suggesting that different contributions to visual analysis for multi-attribute choice might come with different characteristics. As we are focusing on known alternative sets, this space only covers the second half of the decision process axis, where available alternatives are explored and compared. Rather than selective sample points, contributions to assisting decision-making with interactive visualizations

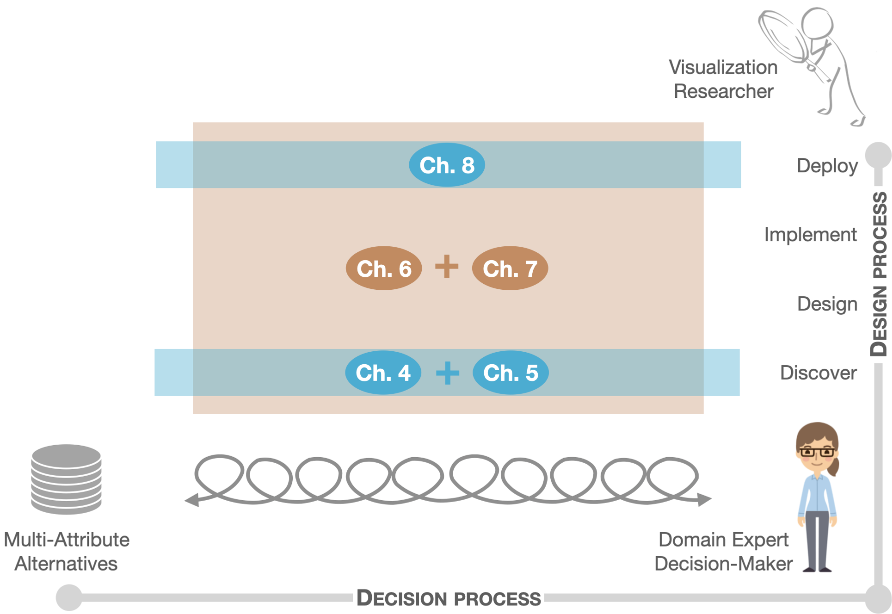

Figure 4.6: Our concept for the visual analysis for multi-attribute choice combines the design process (vertical) and decision process (horizontal) as orthogonal dimensions. Chapters 4, 5, and 8 focus on individual design process stages. Also covering the entire decision process, Chapters 6 and 7 propose complete design studies.

are most likely represented by a path through the space, spanning one or both dimensions. As an example, a contribution might focus on one stage of the design process while being agnostic to the decision process (horizontal path) or it might cover all stages of both the design process as well as the decision process (space-filling area). Note that the axes do not prescribe the exact appearance of specific solutions (e.g., visualization designs, abstractions, etc.). Rather, they are intended to provide a way to think and reason about what problem-driven visualization research has to offer to assist humans in decision processes.

The following chapters instantiate different paths throughout the proposed conceptual space (Figure 4.6). Chapters 4 and 5 improve the characterization of decision problems in the first stage of the design process. Chapters 6 and 7 apply the entire design process to develop visualization artifacts that support experts throughout their decision process. Finally, Chapter 8 assesses the long-term decision support of one of the deployed artifacts to conclude a completed design study.

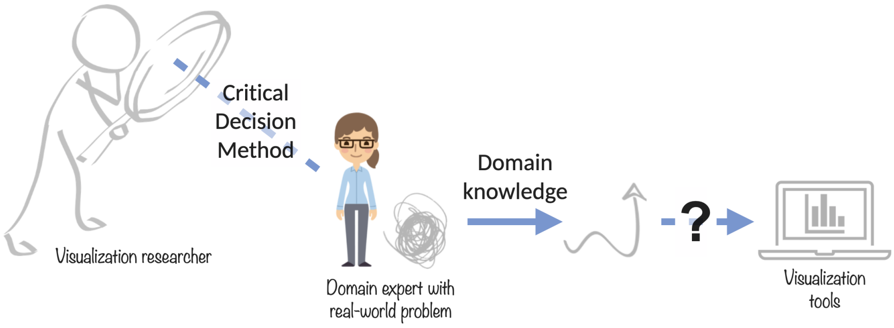

Figure 5.1: This chapter investigates how to reveal the domain expertise and cognitive processes underlying real-world choices.

5 KNOWLEDGE ELICITATION FOR CHOICE TASKS

This chapter investigates a knowledge elicitation method from cognitive science for its ability to facilitate learning about real-world choice tasks in an application domain (Figure 5.1). We are particularly interested in revealing what cues and reasoning strategies are involved and what types of domain knowledge and experience are used. Section 3.1.2 showed that, while domain characterization has become an integral part of visualization design studies, domain expertise can only be acquired from humans through their cognitive processes and methodological prescriptions are rare. We take inspiration from cognitive science, where knowledge elicitation techniques are studied.

In this chapter, we propose one of them, the *Critical Decision Method* (CDM), to the visualization domain in order to contribute to the transition from existing high-level advice towards descriptive steps for domain characterization. Our aims are two-fold: to obtain some initial findings on the use of domain knowledge and experience in choice situations and, equally important, to analyze the applicability of the interview technique for exploring these aspects as part of a task analysis in problem-oriented visualization research. Most parts of this chapter were previously published [51, 97].

What the reader can expect from this chapter:

- A literature review on knowledge elicitation in human science showing that *Cognitive Task Analysis* captures mental processes but has rarely been applied in data analysis settings (Sec. 5.1).
- An introduction to the *Critical Decision Method* procedure as prescriptive steps for eliciting expert judgment (Sec. 5.2).
- A case study on three domain problems suggesting the method's applicability in problem-driven visualization research (Sec. 5.3).

- A reflection on practicability, implications, and directions to prepare the ground for knowledge elicitation to foster decision-focused problem characterization in design studies (Sec. 5.4).

5.1 BACKGROUND ON KNOWLEDGE ELICITATION

What makes a domain expert is the experience and expertise she brings to a particular field of application. This domain knowledge largely consists of concepts, contextual information, typicalities, personal beliefs, learnings, and insights that have been internalized over years of working practice: it is *tacit* knowledge [212].

Different fields of research have evolved around studying tacit knowledge. *Knowledge externalization* [287] aims to convert internalized knowledge to explicit representations, e.g., protocols, that can be reused or shared. Common applications are collaborative sense-making [308] or knowledge-assisted guidance [177]. Articulation is achieved through direct creation of narratives and diagrams, like causal flow charts [301], or indirect inference from user interaction with tools [75]. Among others, externalized knowledge can take the form of labels or annotations [72] or rules from representative decision trees [181]. However, coming from the visualization domain, these approaches either focus on capturing gained knowledge from expert analyses or rely on a user interface to interact with.

Psychologists summarize techniques to capture the unobservable knowledge, mental processes, and goals underlying task performance under the term *Cognitive Task Analysis* (CTA) [178]. Visualization researchers have applied CTA for studying larger groups of domain experts, but not in the context of a design study. Dimara et al. used the critical incident technique to survey the software needs of decision-makers in organizations [68]. Parsons et al. asked participants to retell a past design process to survey the situated knowledge applied by data visualization practitioners [200].

However, it remains unclear how these methods can be applied to domain characterization, because they have been developed for a different context. Knowledge externalization targets tacit knowledge that results from working with data rather than from domain expertise. Cognitive Task Analysis methods center around domain knowledge and naturalistic decision-making but have not been applied in data analysis settings. To make them actionable for domain characterization, CTA methods need to be translated to the domain of visualization research. Beyond a first step in the human factors domain [97], the suitability of CTA techniques for domain characterization has not yet been investigated. For this purpose, we explore the methodological issues of applying the *Critical Decision Method* (CDM) [129] as a representative of CTA.

The Critical Decision Method is a technique to elicit tacit knowledge underlying expert task performance in complex situations [131]. The method grew out of efforts to capture the "knowledge and experience involved in real-world [...] problem-solving" [101]. It has proven useful to investigate dynamic non-routine situations in diverse domains, such as fire fighting or emergency service. Its effectiveness has been demonstrated for a variety of goals, e.g., to develop support systems, design training material, or establish communication strategies [129]. Although these undertakings did not specifically involve visualization, their variability suggests the CDM's applicability also for domain characterization in visualization research. By providing a step-by-step data collection procedure as well as examples for output representations, the CDM carries the potential to address the lack of formal scripts how to conduct domain characterization and how to represent domain characterization results. We consider it particularly valuable for domain problems where the experts' reasoning about data is of particular relevance, like in choice tasks. While the CDM assumes that expertise emerges most clearly during non-routine events, we also see its value to characterize expert reasoning during routine analyses.

The Critical Decision Method uses semi-structured interviews that are often augmented with observations [129]. Thus, it can be considered a variant of "talking with and observing domain experts" [243]. In contrast to previous domain characterization practices, however, it offers more prescriptive guidance. The method aims at a systematic *retrospection* of a situation that involved the participant's expert judgment. The CDM is not meant to replace prospective visualization design. Rather, its retrospective nature is well-suited to understand the vocabulary of the target domain [185] without anticipating future design choices. Traditional approaches to domain characterization ask about current problems or envisioned changes or request users to perform artificial tasks. In contrast, the CDM walks through a past real-world situation, which reveals cognitive aspects of the current problem-solving strategies. It thus helps assess if and how visualization can involve human expertise to better solve the domain problem. The investigated situation, thus, needs to come from the participant's real-world experience. Different targets like cues, knowledge, options, or experience are then probed to understand the expert's reasoning during the situation (see Table 5.1) .

5.2 THE CRITICAL DECISION METHOD PROCEDURE

In this section, we detail the CDM procedure that is used in an interview session to elicit expert knowledge.

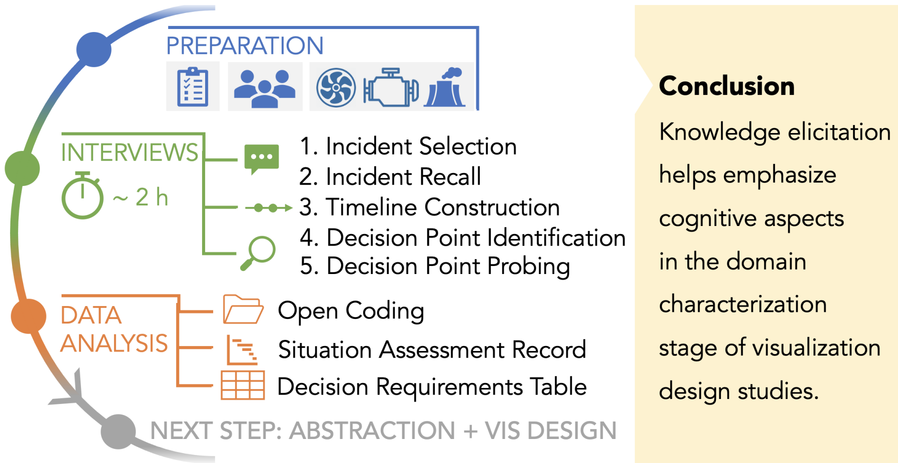

Figure 5.2: Our CDM procedure for domain characterization. To complete a design study, the data analysis results need to be translated into domain-agnostic abstractions to inform the visualization design.

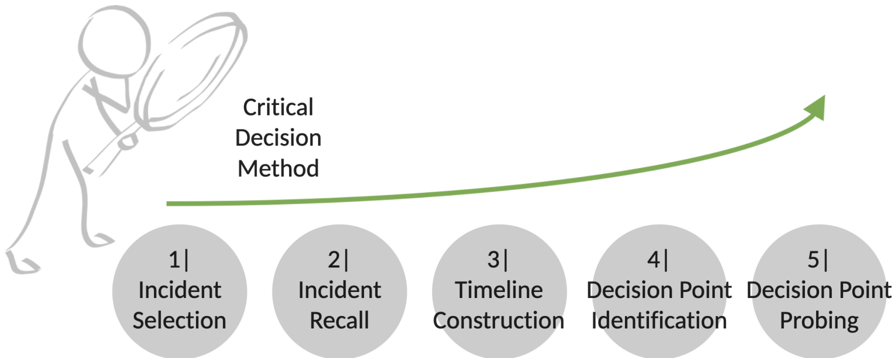

Figure 5.3: Visiting an incident multiple times, the Critical Decision Method gradually elicits the underlying domain knowledge and cognition.

5.2.1 *Five-Stage Data Collection*

Like with other interview techniques, the CDM procedure cannot start until participants have been recruited, ethics approval has been granted, and a questionnaire has been drafted (Figure 5.2, blue). Then, the data collection strategy is to gradually focus on critical cognitive points by sweeping a situation multiple times (Figure 5.2, green) [129]. For this, the CDM procedure provides five stages (Figure 5.3).

1. Incident Selection The goal of the *Incident Selection* stage is to select a task involving competences beyond routine knowledge. The participant should be the primary decision-maker in the situation. To extract true expertise, the task should pose a unique challenge for the participant's competence, i.e., one can expect a difference between the decisions of an expert and those of a novice. It is a pitfall to select a case where participants can rely on formalized procedures.

2. Unstructured Incident Recall The goal of the *Unstructured Incident Recall* stage is to activate the participants' memory and to get a first

Table 5.1: Excerpt of sample questions proposed by Klein et al. [129].

Type	Content
Cues	What were you seeing, hearing, smelling ...?
Knowledge	What information did you use in making this decision, and how was it obtained?
Options	What other courses of action were considered by or available to you?
Experience	What specific training or experience was necessary or helpful [...]?
Basis	How was this option selected [...]? What rule was being followed?
Goals	What were your specific goals at this time?

impression of the scenario. The participant is asked to describe the situation from beginning to end. For example, this might range from loading a data set until an interesting correlation has been found. Interviewers should focus on understanding the story. Interruption for other than minor clarifications is a pitfall.

3. Timeline Construction The goal of the *Timeline Construction* stage is to establish a common understanding among interviewers and participant. From what they heard, the interviewers reconstruct the situation in the form of a timeline. It contains the sequence and duration of events. An event can be an occurrence (like a data point becoming highlighted or a view being switched) or subjective thoughts reported by the participant (e.g., "I would consider this point an outlier"). The timeline is then retold to the participant to identify inconsistencies, add clarifications, and fill in missing details.

4. Decision Point Identification The goal of the *Decision Point Identification* stage is to select relevant decision points in the timeline for a detailed investigation. The interviewers extract those moments where different ways to understand the situation existed or multiple courses of action were possible. Some are obvious from verbal cues (e.g., "I had to decide whether to include this predictor in the selection"). Others involve taking one of multiple courses of action (e.g., looking at one part of the data first), making a judgment that affects the action (e.g., "this shape looks like an anomaly but we can safely ignore it"), or making a choice that a novice might have made differently. The granularity of decision points can be adapted.

5. Decision Point Probing The goal of the *Decision Point Probing* stage is to better understand the meaning of information for the participant's assessment of the situation. The interviewers work through the decision points and ask for elaboration. Different probes can be applied for this purpose based on the interviewers' research interest.

Table 5.1 lists example questions about how cues, prior knowledge, or different options influenced the participant's course of action.

Previous stages can be revisited to gradually refine the timeline and decision points. An interview is expected to last about two hours. This can vary depending on the application (e.g., replace timeline construction by observation or conduct decision point probing during breaks). Klein et al. recommend to share the interviewing responsibilities among two interviewers and to record the sessions [129].

5.2.2 *Data Analysis and Output*

The Critical Decision Method does not prescribe a data analysis approach (Figure 5.2, orange) because it depends on the research questions motivating the undertaking. In general, coding is used to prepare the ground for converting the interview data to different representations that describe domain knowledge, reasoning, and task activity [101]. Hoffman et al. recommend to tag cognitive functionality like perceptual cues, decision points, and situation assessments [101]. Klein et al. present representations that worked well for their applications [129]. We highlight two of these artifacts that we found to particularly match the purpose of domain characterization.

A *situation assessment record* (SAR) reflects the expert's understanding of the dynamic evolution of an incident. It specifies the turning points of a situation together with underlying cues, experience, knowledge, goals, and actions. Klein et al. propose different formats for SARs [129]. Table 5.2 shows an SAR for water turbine design: an existing turbine is analyzed to derive potential directions for improved running behavior. The granularity of entries can be adapted to the researchers' needs. New events or insights cause the expert to abandon prior goals and prioritize new goals (shift). For example, identifying the cavitation on the turbine blade as the major problem changes the engineer's goal from analyzing the existing turbine towards optimizing the blade geometry (Table 5.2, SA 3). Sometimes the goals are maintained but new information enhances what was originally known (elaboration). For example, the cavitation happens at the leading edge of the blade (Table 5.2, SA 4).

A *decision requirements table* contains details on the judgments that were involved in performing the observed task. The columns specify what particular decisions were made, why they were difficult to make, how they were made, and what supporting information was used. The rows correspond to the decision points identified in the situation assessment record. In this way, the decision requirements complement the experience, goals, etc. in the SAR. For example, the difficulty to investigate multiple operating points (Table 5.3, 2nd row) effectively extends the description of SA 2 in the SAR. The prescribed structure of Table 5.3 allows for a comparison even across situations.

Table 5.2: Situation assessment record (SAR) for water turbine design.

Situation Assessment 1	Plausibility check
Cues	Deviation from onsite or testbed measurements
Goals	Validate simulation model
Decision Point 1	If invalid, calibrate, otherwise proceed
Situation Assessment 2	Analysis of current setup (shift)
Cues	Pressure distribution on blades (heatmap), performance indicators (e.g., torque and power output)
Knowledge	Challenges/trade-offs associated with design
Experience	Problems in previous operation of the turbine (e.g., cavitation at leading edge)
Goals	Understand strengths/weaknesses of existing turbine, identify potential directions of improvement
Decision Point 2	Address cavitation on current turbine blades while keeping its strengths
Situation Assessment 3	Optimize blade geometry (shift)
Cues	High curvature in pressure line
Knowledge	Correlates with high blade angle change
Experience	Avoid by shifting camber towards leading edge
Options	Change blade angles or meridional length
Goals	Achieve constant pressure change
Decision Point 3	Shift leading edge towards inlet
Situation Assessment 4	Further optimization (elaboration)
Cues	Pressure/velocity distribution (heatmap, streamlines, sweeping plane), performance indicators, pressure/angle across blade length (line chart)
Basis	Trial-and-error exploration of design space, per step analyze what went right and how to improve
Knowledge	Flow behavior, how parameters relate to side effects, constraints (e.g., construction volume)
Experience	Dependencies known from previous projects, operating permit requires trading 2-3% less efficiency for fish-friendliness, operating conditions might change during the project
Goals	Understand how geometry affects water flow, trade off efficiency and fish-friendliness
Decision Point 4	Proceed with most preferred turbine design
Situation Assessment 5	Improvement potential (shift)
Cues	Efficiency curves new design vs. existing turbine
Experience	Desired flow and pressure lead to high efficiency
Goals	Predict savings/earnings for customer
Decision Point 5	Implement chosen design

Table 5.3: Decision requirements table for water turbine design. The rows correspond to the decision points identified in the SAR (Table 5.2).

	What is the decision?	Why is it difficult?	How is it made?	What is the aid? How does it help?
1	Determine plausibility of simulation results	Simulation model uncertainties, inaccurate real-world measurements, tolerable deviation not known a priori	Compare real-world measurements to simulated values	n.a.
2	Analyze strengths and weaknesses of current turbine	Only one operating point can be investigated at once	Model and simulate the existing turbine and investigate its performance	3D visualization of pressure/velocity/etc. distribution across geometry hints at inconsistencies
3	Optimize blade geometry	Non-linear relationship between geometry and water flow	Identify non-optimal pressure distribution across blade, change blade geometry, and observe effect	2D line charts that convey the pressure distribution and blade angle across the blade length, high curvature hints at non-optimal geometry
4	Choose most preferred turbine design	Performance increase in one operating point might come with reduced performance in another operating point, manual exploration of design space, intuition difficult to formalize	Trial-and-error approach, iterative manual changes of the geometry and direct observation of performance changes	2D line charts and 3D visualization of output parameters convey distribution across geometry, color encoding draws attention to peaks, instant feedback helps see the relationship between geometry changes and flow changes, knowledge of turbine designer hints at direction to improve
5	Assess improvement potential of chosen option	Affected by entire operating range, hard to predict savings	Compare choice to existing design, critical cue is efficiency	Superimposed efficiency curves of both designs enable direct comparison

5.3 APPLYING THE CRITICAL DECISION METHOD IN PRACTICE

This section will present the CDM steps that we followed in conducting three domain characterizations to understand how engineers approach optimization problems that are based on simulation data. Some aspects of the CDM study have already been published in a conference paper [97]; however, the data analysis covered a different purpose. The previous work investigated the feasibility of the CDM from a human factors point of view. Here, we reflect on the methodology from a visualization design study perspective. Besides a summary of the elicitation procedure, we discuss the method's applicability for the purpose of domain characterization.

5.3.1 *Application Background*

We studied the domain practices in three different applications from the field of engineering design: optimizing a water turbine, an electric drive, and the operation modes of a power plant. The expectations regarding engineered systems are constantly rising: customers ask for high-quality products that are available at little cost and in a short time frame. Environmental requirements play another important role. As such, all three applications deal with multi-attribute decision-making, a core goal of visualization [67]. It is challenging because rationality is often complemented with intuition [128, 200]. We were interested in the experts' mental processes and domain knowledge involved with trading off multiple criteria when choosing the most preferred design and operation mode of the system. The gained insights might inform the design of a visualization that supports the experts in navigating the design space and applying their expertise for trade-off strategies.

In the following, we provide context for each application. All three optimization scenarios involved interactive visualizations previously grown in the target domain. Note that our main purpose was not to evaluate the visualization designs but to assess in what sense they do (not) help achieve the experts' goal.

Water Turbine Design A water turbine converts water flow into electric energy. In the context of modernization, an existing turbine is to be optimized regarding its running behavior given dynamic operating conditions like water throughput. A typical problem is cavitation on the turbine blades that is caused by low water pressure and can lead to serious mechanical damage. The optimization is characterised by repeated geometry changes followed by an exploration of their effects until the engineer is satisfied. This process largely relies on the experience of the designer regarding how the turbine geometry interacts with the water flow. A typical criteria conflict faced by turbine designers is that between turbine efficiency and fish-friendliness.

Power Plant Optimization A thermal power plant burns fuel to convert the heat energy into electric energy. Changing the fuel type without further adjustments might lead to fuel remaining unburned. Operators thus perform an iterative design space exploration to find an operation mode (e.g., input temperature, valve and damper positions) that reduces the amount of unburned fuel while maintaining operational characteristics like nitrogen oxide emissions and exit temperature. A challenge are coupled physical effects such that eliminating one problem might cause an unexpected problem elsewhere.
Electric Drive Design Electric drives convert electric energy into linear or rotational motion. They are at the core of many applications, from electric toothbrushes over ship propulsion to transportation systems. Optimizing electric drives means to specify their geometry, material, winding patterns, etc. such that their performance optimizes given requirements like cost-efficiency, durability, or construction volume. For this, the operational behaviors of many different electric drive designs are simulated and genetically optimized. This gives a set of objectively equal solutions, where no criterion can be further improved without sacrificing another criterion. Among the resulting set of solutions, it is the responsibility of the design engineer to choose the most preferred compromise. This is challenging, because the number of Pareto-optimal solutions is often quite large and optimization criteria are typically conflicting. A detailed description of this use case can be found in Chapter 6.

5.3.2 *Implementation of the CDM Procedure*

In this section, we detail our realization of the CDM. Explanations applying to all three studies are accumulated. Where relevant, we explicitly differentiate the studies.

Following the recommendation by Klein et al. [129], we recorded a 90-minute remote session for each application. A pair of elicitors, one visualization expert and one human factors expert, shared the responsibilities for asking questions, taking notes, and analyzing the protocols afterwards. Due to the spatial distance of the elicitors and participants, the interviews were conducted remotely. Ethics approval was obtained from the Faculty of Engineering at the University of Nottingham prior to conducting the studies. The participants' involvement in the study was voluntary and we offered a compensation for their time. They were free to withdraw at any time. In case of withdrawal, we did not keep any of their data. A participant information sheet including these information was distributed to the participants prior to the interviews.

In preparation for the semi-structured interviews, we compiled a set of open questions to give us a general overview of important decision characteristics to cover. The questions were grouped according to the stages in the CDM procedure. We used these probes as an aid to keep our elicitation focus, while also probing upcoming aspects on the fly. We began each session by introducing its goal: to investigate the cognitive aspects involved in trading off multiple criteria with the help of interactive visualizations. We then collected the demographics and asked the participants about their general way of dealing with design optimization in their daily work. From there, we worked through the CDM stages as described below. Together with the participant information sheet, the complete questionnaire can be found in the supplemental materials[1] of our previous publication [51].

1. Incident Selection All three applications required expertise beyond the general routine knowledge of a competent individual. The participants were full-time employed engineers and performed multi-attribute decision-making with visualization support as part of their daily work. They had ten to fifteen years experience in engineering design. This qualified them as experts [129]. An obvious but critical prerequisite was that participants were willing to share material related to the optimization during the interview. We completed the incident selection stage via e-mail in advance to spend the interview time on the actual knowledge elicitation.

2. Unstructured Incident Recall As specified in the original method [129], we requested the participants in all three studies to verbally provide a brief overview of their application. We did not interrupt them and focused on understanding the story. The recounts covered the purpose of the engineered system to be designed (e.g., convert water flow into electric energy), the parameters of a design option (e.g., geometry and water throughput as design parameters and efficiency, power output, fish-friendliness, and mechanical robustness as objectives), and the approach taken to arrive at a preferred design. All three applications involved four to six optimization criteria whose coupled effects raise the challenge of improving on one criterion without reducing the performance in another criterion. Generally, we observed two different types of approaches: 1) an iterative trial-and-error exploration with alternating parameter variation and inspection of the results and 2) a genetic optimization approach with automatic variation of parameter settings to compute hundreds of solution candidates to inspect. Going beyond the original CDM method, we classified the reported approaches as a priori, a posteriori, or interactive optimization [174], depending on when in the optimization process the participants articulated their preferences. This helped us anticipate the chronology of the timeline in the next stage. In the water turbine study, the expert made use of her expertise in each iteration of exploring the effects of geometry modification. Thus, it belonged to the interactive methods.

3. Timeline Construction + 4. Decision Point Identification The CDM procedure provides that both stages are performed in parallel [129].

1 https://ieeexplore.ieee.org/abstract/document/9990992/media#media

In all three studies, we reconstructed the process of deciding for a preferred design option by having the participants walk us through material related to their applications in sequential order. This slightly deviated from the original method, where the timeline construction is based on the *Unstructured Incident Recall*. The timelines were similar across all three studies. They included an initial plausibility check of the simulation model, an iterative exploration and optimization phase, and a confirmation of the final choice. The water turbine design and power plant optimization started off from an existing case that served as a reference to be improved. The electric drive design focused on applying customer specifications and technical expertise to select the most preferred compromise from a large set of candidates. Table 5.2 shows the evolution of the water turbine design. It included modeling an existing turbine (SA 1), analyzing its performance (SA 2), iteratively modifying the blade geometry and observing its effects (SA 3 + SA 4), and comparing the optimized geometry to the initial one (SA 5). The decision points are also highlighted. Some were obvious from verbalization, e.g., *"now we [...] want to start the simulation"* or *"I can now change the operating point"*. Others included subjective assessments of the simulation results, e.g., *"I see that the flow below the runner is good"*. While the CDM recommends to capture both the sequence and duration of events, we omitted the duration in all three studies, because the optimizations were not time-critical.

5. Decision Point Probing In line with the CDM method, we asked for additional details on some decision points. We selected those turning points where the participants seemed to have particularly relied on their expertise to steer the optimization direction. Which decision points are selected for deeper investigation and what details are carved out depends on the purpose of the study. In all studies, we browsed the sample questions [129] for inspiration (see also Table 5.1). Given our focus on choosing the most preferred design, we particularly asked about available options and how an option was selected, e.g., "Based on what constraints did you exclude these options?". We also probed for visual cues, e.g., "Where do you see that in the visualization?", and prior knowledge, e.g., "How do you know what parameter to adjust next?", that helped the experts gain insights at each optimization step. For the water turbine study, Table 5.2 contains the utilized question types for each decision point.

Data Analysis and Output Upon completion of all three interview sessions, we analyzed our protocols and recordings. While the CDM does not prescribe a concrete analysis method, it recommends any form of coding to get started [101]. In the *Decision Point Identification* and *Decision Point Probing* stages, we already tagged the decision points with their underlying cues, experience, goals, etc. After all sessions had been completed, we used a qualitative coding methodology [260] to identify meaningful recurring topics in the data. Given our study purpose, we particularly watched out for common goals, needs, approaches, and difficulties that emerged from the responses of the participants. Still, we followed an inductive, data-driven approach that did not assume a set of pre-established codes. From the protocols, recordings, and code set, each interviewer further derived two of the proposed CDM outputs [129] per session: a situation assessment record describing the situation as a series of decision points and a decision requirements table specifying the what, how, why, and aid of the particular decision points. We exemplary depict the situation assessment record (Table 5.2) and decision requirements table (Table 5.3) for the water turbine study. The results for the remaining two use cases can be found in Appendix A. We performed the coding and artifact generation independently. We then discussed the results and jointly refined the codes and artifact representations. The resulting themes included aspects ranging from intuition and subjective judgment over stakeholders to post-decision workflows.

5.3.3 *Practical Considerations*

This section presents the methodological issues we observed when applying the CDM for the purpose of domain characterization for real-world decision problems. We illustrate method properties we recognized, domain characteristics that the method helped elicit, and things that worked (less) well.

Capturing the cognitive aspects associated with real-world decision-making requires a careful consideration of the decision-maker's subjective conditions as well as environmental conditions (e.g., the complexity of options to consider or time pressure). This involves characteristics and dependencies that cannot be easily quantified. Thus, a qualitative method like the CDM is beneficial to elicit and abstract meaningful information that can inform visualization designs across application domains.

However, qualitative methods are time-consuming and the CDM is no exception. The 90-minute sessions with the experts were followed by a time-intensive data analysis by us as interviewers. In design studies, however, the outcome of the elicitation is critical for all subsequent layers [185]. Given the contextual richness of the CDM responses, we thus consider the efforts justified. Our impression is that domain problems that are expected to be cognition-sensitive and could only be described by a series of low-level tasks benefit most from a characterization using the CDM.

We originally performed another CDM session on a signal filter optimization. It failed, because the optimization relied on a routine weighting strategy rather than expertise. We did not realize this until the *Unstructured Incident Recall* stage. The participant was also in an early PhD stage and thus did not have enough expertise. We eventually decided to discard this session from our collection.

In each of the remaining studies, the *Unstructured Incident Recall* stage conveyed an initial idea of the domain problem as intended. The participants could give an overview of their situations on an appropriate level of abstraction, which kept the entry barrier for the subsequent stages low without too many details. Where needed, we adjusted the level of abstraction by asking clarifying questions.

The *Timeline Construction* walk-through, in contrast, contained a lot of meaningful details. Remembering a past situation in detail is difficult and may provoke a mismatch between user recollections and their actual actions [243]. While we did not evaluate this, our participants did not seem to have difficulties with remembering the past incidents. It might have helped their memory that we, aside from the original CDM method, encouraged them to bring documentation material. The water turbine study, in particular, involved technical details, e.g., *"this is a hand-coded mesh generator"*, that sometimes distracted us from the actual reasoning process.

The CDM centers around an incident from the real-world practice of experts. In contrast to other approaches, where the participants might be confronted with unfamiliar tasks, the CDM allows the experts, not the interviewers, to choose a meaningful situation that they know well and feel comfortable with as the subject of investigation. This enables participants to provide a comprehensive review of the situation and their reports are not biased by unfamiliarity with the task. All of this helps avoid the domain threat of mischaracterizing the problem (cf. the top level of the nested model [185]). An immediate validation of that threat is naturally incorporated in the *Timeline Construction* stage of the CDM procedure: retelling the constructed timeline to the participant. Deviating from the original procedure, we skipped this step, thus missing the chance to validate the constructed timeline and decision points during the interview. In retrospect, we should have followed the original procedure or included a post-interview validation of our results.

How much the CDM can teach researchers about a target domain might depend on the knowledge gap commonly associated with application-driven visualization research [296]. Considering an incident from the participants' experience leaves the interviewers fairly naïve about the discussed scenarios. It is not necessarily a drawback: interviewers who know little about a domain tend to probe more. Still, we underestimated the mental effort for decision point identification and probing without much prior knowledge about an incident. In the *Incident Selection* stage, we only requested as much information as needed to assess a participant's suitability for the study. We found it difficult to reconstruct the timeline and choose appropriate probes on the fly during the interview session. Thus, we largely relied on questions that we prepared beforehand without knowing the incidents in detail. We also put together the sequence of decision points in retrospective, i.e., upon completion of all three sessions. In retrospect, once the participants were found to be suitable, we should have asked for further material to familiarize ourselves with the particular incident prior to the interviews. While this might help, it still remains an open question how to succeed in spontaneous timeline construction.

Our independent data analysis results showed a broad consensus regarding the content, especially with respect to those moments of an observed situation that we considered decision points. This concurs with Klein et al., who found that inter-observer variability refers to the significance of a decision point rather than its presence [129]. For a subsequent task abstraction this suggests that disagreements between visualization researchers might mainly relate to *why* a task is performed [38]. Although both interviewers based their analysis on the same situation assessment example, we found the resulting records to significantly deviate in their format, i.e. the mapping between decision points and situation assessments as well as their granularity. These deviations propagated to the decision requirements table, because we transferred the decision points from the situation assessment record to the table rows. Explicitly agreeing on a template beforehand might further reduce the risk of discrepancies.

The systematic CDM procedure revealed what types of domain knowledge and expertise the users carry, e.g., *"the designer knows about the parameter options and side effects"* or *"operators usually trade 2% to 3% less efficiency for fish-friendliness"*. We further learned what cues steer their attention, e.g., *"co-occurrence of high oxygen and high temperature"* or *"too much curvature in the pressure lines"* (compare Table 5.2, SA 3, Cues). The CDM also made explicit how the experts' goals varied with the situation focus, e.g., from understanding the status quo over reducing unburned fuel to maintaining a reasonable cost-benefit ratio. To conclude, by revealing the role of user expertise in task performance, the CDM has the potential to effectively foster the *appropriateness* of a visualization [176], i.e., its benefit for supporting a given task.

5.4 IMPLICATIONS FOR VISUALIZATION RESEARCH

This section reflects on our experiences with the Critical Decision Method more generically against the background of visualization research. For this purpose, we indicate how the method aligns with existing approaches from visualization research. We also highlight questions that remain unanswered and how future research can address these open issues.

5.4.1 *Reflections on Domain Characterization*

Explicitly describing tacit knowledge has been an ongoing challenge. We experienced that the CDM comes quite close by producing artefacts that explicitly describe applied expertise, subjective judgments, and contextual effects. This is particularly relevant for visualizations, because they are highly dependent on the goal, task, and context of their usage. The decision points in the situation assessment record (Table 5.2) translate to individual contexts that might require different design choices. The extracted goals hint at what the participants want to do (rather than see in a visualization), i.e., the analysis tasks they need to work through. The timeline might inspire narratives for downstream validation of a visualization. For example, in a field experiment [45], where realism is manipulated by asking participants to perform specific tasks, the timeline can frame a particular setting.

We experienced that probes like "what rules did you follow to make this decision?", or "what were your specific goals at this time?" helped experts concentrate on what they want to do rather than what a visualization solution might look like. On the other hand, we identified probes for perceptual cues, i.e., "what were you seeing?" or "what caught your attention?", as one possible starting point for deriving design requirements. They revealed characteristics in the data to be emphasized by a future visualization design. Answers to the role of (visual) aids in the decision requirements table (Table 5.3) also pointed towards potential entry points for visualization support. Similarly, responses to experience and knowledge probes could help design interaction techniques. In combination with perceptual cues, they might also inform the integration of guidance into a visualization system [46]. We did not use probes for hypotheticals (e.g., "what difference would it make if ...?") ourselves but we expect them to help anticipate the consequences of different design choices and raise particular awareness for potential pitfalls.

In prior projects where we did not apply the CDM, we conducted think-aloud walkthroughs where we asked experts about a visualization's rather general context of use. In contrast, the comprehensive, yet systematic, procedure of the CDM helped us stay focused on decision points that notably revealed expert knowledge. Its output can pave the way for turning incidents into abstractions and subsequent design choices. More precisely, the CDM helped us identify and describe critical decision points that can be used in subsequent layers of the design process.

Rather than a strict recipe, the CDM can be seen as a framework where implementations can be chosen according to the research objective. It is open to being combined with dedicated requirements engineering techniques from visualization research. Up to now, domain characterization focuses on tasks [185], but less on domain knowledge elicitation. Complementing existing approaches with the CDM allows for a complete picture of the target domain, including both the task-based and the knowledge-based perspectives.

For example, the CDM is one of many ways to implement the *discover* stage of the nine-stage framework [243]. It contributes to preventing pitfalls PF-15 ("ignoring successful aspects") and PF-17 ("focusing on visualization solutions"). The *Unstructured Incident Recall* with passive interviewers provokes PF-16 ("expecting talking and passive observation alone to work") at first sight, but additional think-aloud sweeps of the incident compensate this. As a qualitative approach, the CDM also seems to provoke PF-5 ("insufficient time from collaborators"), but it actually makes efficient use of expert time, leaving the time-consuming part to the researchers.

The CDM builds upon a holistic consideration of probe types like cues, goals, or knowledge. They convey a comprehensive picture of cognitive turning points in a domain problem. Although targeted at decision points in the first place, insights gained through these probe types also hold the potential to advance existing task descriptions. Furthermore, in analogy to the multi-level typology of abstract visualization tasks [38], the *what*, *how*, and *why* classification in the decision requirements table (Table 5.3) can inform a visualization-oriented characterization scheme for decision points. The CDM can then serve as a systematic data collection method to inform the creation of a taxonomy of decision tasks. In a similar way, it could help identify different domain knowledge types and their representations to inform endeavors in knowledge-assisted visualization.

To conclude, any domain characterization technique will highlight some aspects of the problem domain and de-emphasize others. Depending on the research objectives, multiple methods can be combined to arrive at a concise understanding of a target domain. In this sense, the CDM is a valuable addition to the portfolio. Its output in the form of open coding, situation assessment record and decision requirements table provides a good basis for a) discussions and reflections among visualization researchers and domain experts and b) the subsequent definition of abstract tasks, requirements, and mental models to inform the visualization and interaction design.

5.4.2 *Current Limitations and Future Research Paths*

This work is a first step towards integrating knowledge elicitation from cognitive science into visualization research. Further investigation is needed to back up our experiences and turn the *Critical Decision Method* into an actionable model for visualization researchers. We discuss current limitations and highlight the future research paths that emerge from our work.

A major limitation is that, while we applied the interview technique on real-world use cases for domain characterization activities (Section 5.3), it has not been embedded in a complete application of the design study methodology. We need to further investigate how the CDM can be concretely applied to problem-driven visualization design. This includes an adaptation of the terminology (e.g., from *incident* to *analysis*), particular probing of aspects related to analysis (e.g., data quality or correlations) and visualization (e.g., correlation or visual representations), and more concrete dedication of each CDM step to the goals of domain characterization. To arrive at a complete domain characterization model, the CDM procedure also needs to be linked more tightly to the subsequent layers of visualization design. Up to now, the technique's findings have not been used as a basis for the design of visualization tools. While we reflected on its conceptual interfaces to other methods that are commonly used in design studies (Section 5.4.1), practical experience with the integration in a design study workflow is not yet available.

The previous point is also connected to finding the best approach to evaluate the effectiveness of the CDM framework. Among others, Marai and Möller propose *significance* and *pragmatic adequacy* as evaluation criteria for theoretic contributions to visualization research [159]. We hope to have motivated the significance of the CDM to the visualization field. Yet, it is to be confirmed, whether the expert knowledge elicited by the interview technique indeed contributes to better visualization practice, i.e., improved results of the subsequent visualization design layers. Reporting on a complete design study clarifies whether the knowledge elicited by the CDM actually helps improve the results of subsequent visualization design layers. This might also target professional practitioners as opposed to design studies conducted in visualization research [199]. With a number of observations collected, we might also be able to identify meaningful practices that can serve as guidelines on how to design visualizations that foster the exploitation of human cognition and knowledge for analytic tasks. Validating practical experiences with the CDM should also include data analyses with different high-level tasks. This can even be extended to a set of real-world case studies from different domains that discuss the CDM procedure step by step and compare it to alternative approaches. By replicating domain characterizations from existing design study papers, the previous approaches could be compared to a CDM domain characterization, highlighting the differences that stem from using the CDM. We note that this would mean to compare empirical methodologies rather than techniques, which is not commonly done in the visualization community.

5.5 CONCLUSION

This chapter examined the feasibility and significance of the Critical Decision Method for eliciting expert knowledge involved with professional decision-making. To develop visualisation techniques that are effective, we need to better understand the cognitive processes that are at the core of decision-making with visualizations [197]. This chapter considered ways of capturing such cognitive aspects by reviewing knowledge elicitation methods that have proven successful in cognitive science, but have not yet received much attention in visualization research. One of them, an interview technique called *Critical Decision Method* (CDM), was found to combine multiple properties that are relevant for learning about a target domain in visualization design studies. It employs a systematic five-stage procedure that revisits a domain problem multiple times to gradually elicit expert knowledge. Despite its significant effort for data acquisition and analysis, applying the procedure to study three applications from engineering design showed its potential to compile rich collections of cues, subjective judgments, and contextual effects involved in decision-making.

The CDM provides an alternative way of learning about domain experts and the conditions that frame their task performance. Its focus on real-world incidents aligns well with the required realism in tasks, data, and users for understanding work practices [138]. The CDM particularly encourages participants to reflect on their own cognitive processes. It suggests a novel perspective on domain characterization by favoring decision points over tasks. While it is grounded in the context of professional decision-making, we also see its potential for casual decision-making. To summarize, we found it to be a promising way to emphasize cognitive aspects in decision-making scenarios.

Our experience is a first hint towards general suitability of the method. However, further investigation is needed to arrive at an actionable method to inform problem-driven visualization design. This includes a more concrete dedication of the procedure to the needs of visualization research as well as studies to evaluate its effectiveness. From there, we could derive methodological guidelines that guide the consideration of cognitive aspects identified by the CDM to improve decision support visualizations.

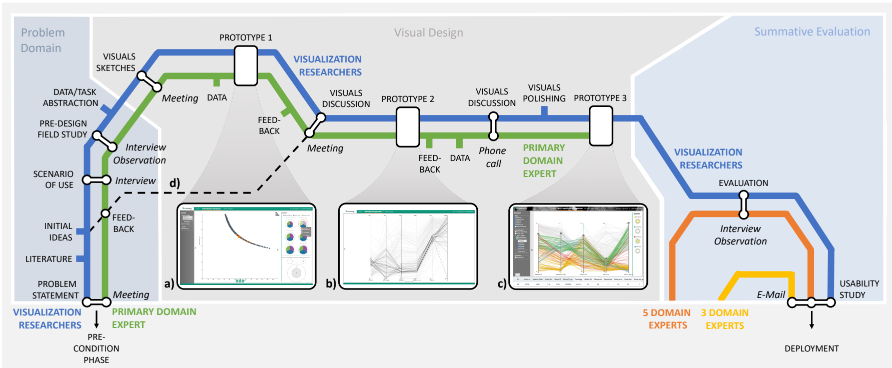

Figure 6.1: An overview of our design process, inspired by storyline visualizations [264]. A mechatronics scientist (green) accompanied us (blue) during the domain characterization and visual design stage. The process involved two intermediate prototypes (a, b). The final prototype (c) was qualitatively 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.

6 PAVED: PARETO FRONT VISUALIZATION IN ENGINEERING DESIGN

This chapter presents a design study on the role of multivariate visualization to support multi-attribute choices in engineering design applications.

Chapter 3 raised the need to study the effectiveness of visualizations on real-world multi-attribute choices made by real users. It showed that, although design studies have become increasingly popular in visualization research, no design study targeting choice over analytic tasks has been presented yet (Section 3.1.1). This calls for a systematic visualization design study on multi-attribute choice. Section 2.4 identified the potential of visualizations to support choices in engineering design but reveals a lack of dedicated visualization research.

This chapter describes a 1.5-year collaboration with mechatronic scientists on supporting the design and optimization of electric motors with visualizations. The engineers' choice task is to select the most preferred electric motor design from a large set of Pareto optimal solutions (for details on Pareto optimality, we refer to Section 2.2.2).

A primary goal in this chapter is to accurately characterize the domain-specific data and choice task and to develop a validated visualization design that effectively supports the engineers in solving their task. Familiarizing with the design and optimization of electric motors lays the foundation for an abstraction of the involved data and tasks, ultimately leading to a set of design requirements. These in turn motivate our visual design decisions, which we implement and validate in a case study with motor designers.

A secondary goal in this chapter is to share our experience with the user-centered design process, in particular regarding the collaborative aspects. We took methodological inspiration from Sedlmair et al.'s nine-stage framework [243] and Munzner's nested model [185]. A number of design studies have already reflected on their implementation, but we particularly view our findings in the context of the optimization of electric motors as an instantiation of multi-attribute choice. This includes repeated observations of the mechatronics engineers' expectations and acceptance regarding the visualization support.

From two different perspectives, these goals contribute to a better understanding of how multi-attribute choices in real-world applications can be effectively supported with visualizations. Our lessons learned offer guidance to other visualization researchers targeting choice problems in engineering design or alternative domains. The design study presented in this chapter has been previously published in a scientific journal [54] and a technical journal [53].

What the reader can expect from this chapter:

- The first design study dedicated to multi-attribute choice showing the potential of working closely with real expert decision-makers to support their ill-defined decision problems.
- A careful characterization and abstraction of the data, tasks, and requirements related to decisions in drive design (Sec. 6.1).
- A description of the collaboration and design process (Sec. 6.2).
- Four design rationales for visually assisting decision-making with the potential to generalize across applications (Sec. 6.3).
- The design of *PAVED*, an interactive parallel coordinates visualization that offers a compact overview of alternatives and simple interactions for incremental preference construction (Sec. 6.3).
- An observational study to assess the tool's domain usefulness showing that PAVED helps decision-makers learn what level of performance is achievable under different conditions (Sec. 6.4).
- Reflections on visualization needs in electric drive design as well as visualization design and methodological guidelines to provide orientation for other visualization researchers (Sec. 6.5).

6.1 DOMAIN CHARACTERIZATION AND ABSTRACTION

Section 2.4.1 described how multi-attribute choice generally manifests in applications within the field of engineering design. As one of these applications, we develop a detailed characterization of the targeted problem regarding the design and optimization of electric motors. Based on an introduction to the design of electric motors (Section

6.1.1), we provide an abstraction of the data and tasks that engineers face when searching for the most preferred design (Sections 6.1.2 and 6.1.3). From these abstractions, we derive design requirements to be addressed by the visual design (Section 6.1.4).

6.1.1 *Background on Electric Motor Design*

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 [306]. 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 [70]. 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 [306]. 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* [248] (formerly *MagOpt* [247]). 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

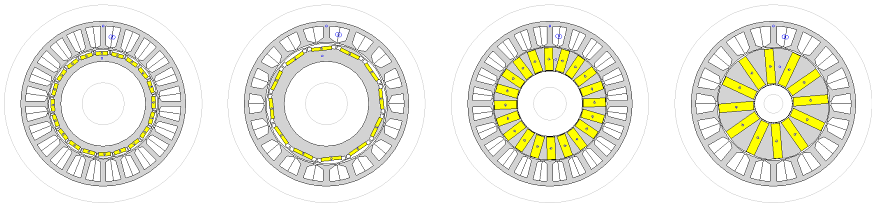

(a) Neodymium, 20 magnets (b) Neodymium, 12 magnets (c) Ferrite, 20 magnets (d) Ferrite, 12 magnets

Figure 6.2: 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.

with an efficient drill down, perception of redundancies and conflicts among criteria, and sensitivity analysis.

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.

6.1.2 *Data Abstraction*

We first provide information on the simulation and optimization approach that is used by the domain experts to generate the set of motor designs to choose from. We then move from the domain-specific details to the data abstraction.

PMSM Simulation Model

The simulation considered as a running example throughout this work describes the operational behavior of a permanent magnet synchronous motor (PMSM). PMSMs are increasingly used for applications where high efficiency and power density are crucial. They can be found in the automotive sector, home appliances, or medical devices but continue to expand into all other areas of use of electric motors. In this case, the motor's 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 6.2). 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.

Abstraction

Simulation models are basically input-output models that approximate a function $X \rightarrow Y$ mapping some input dimensions $X = \{X_1, ..., X_n\}$ to a number of output dimensions $Y = \{Y_1, ..., Y_m\}$. In line with the terminology introduced in Section 2.1.6, 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 metadata. The union $(\mathbf{x}, \mathbf{y})$ of a design option $\mathbf{x} = (x_1, ..., x_n); x_i \in X_i$ and its performance $\mathbf{y} = (y_1, ..., y_m); y_i \in Y_i$ as provided by a simulation run is called *alternative*.

Section 2.4.1 highlighted that the challenge of engineering design lies in the absence of a direct inverse relation $Y \rightarrow X$. 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.

6.1.3 *Task Abstraction*

As Ullman states: *"[engineering] design is decision-making"* [281]. Section 2.4.1 discusses characteristic properties of the engineering design domain. To summarize, the engineer's 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 in Section 4.3.2.

A variety of strategies for multi-attribute choice have been studied in decision theory (Section 2.1.4). In their position paper, Torsney-

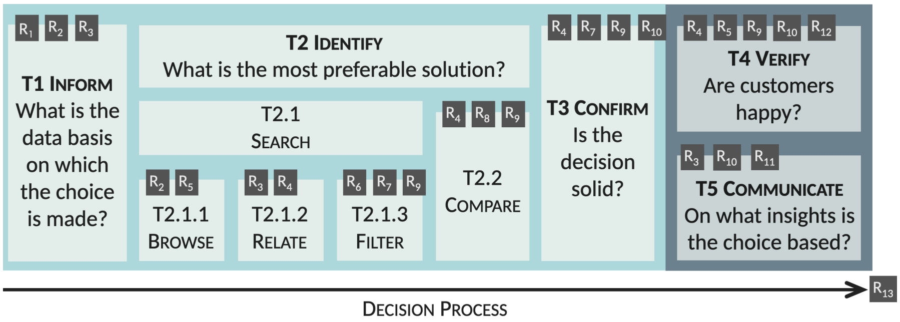

Figure 6.3: An abstraction of the analysis tasks for multi-attribute choice in engineering design. T1 to T3 reflect the tasks that the engineer faces. T4 and T5 involve the customer as additional stakeholder. Derived design requirements R_1 to R_{13} are assigned to the tasks.

Weir and colleagues propose to consider such strategies for visual tool design [269]. We therefore began the task characterization by classifying the decision strategy of our primary domain expert. According to the framework described by Payne et al. [203], 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 6.3) is based on selected tasks from 11 taxonomies that we found in the visualization literature. Details on the task extraction and analysis questions underlying the selected tasks can be found in Appendix B. Although our task abstraction aligns with commonly employed engineering design steps (Section 2.4.1), it is not specific to this 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 [281]. 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. In working through these subtasks, the decision-maker basically conducts a trade-off analysis (Section 2.2.3).

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.

Comparing these tasks to Simon's decision process model (Section 2.1.4), we notice that our identified tasks loosely map to the stages he proposes: inform (T1) maps to the intelligence activity, identify (T2) to the choice activity, and confirm (T3) to the review activity. The communication need we identified (T4 and T5), however, does not appear as such in Simon's model.

6.1.4 *Design Requirements*

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

- **R1 Validation** – Show criteria ranges for simulation steering
- **R2 Overview** – Provide an overview of complete options such that any attribute value of any option can be retrieved
- **R3 Criteria Relations** – Highlight redundant or conflicting criteria
- **R4 Trade-offs** – Support subjective perception of superiority
- **R5 Filter** – Support perception of the effect of constraints
- **R6 User Interaction** – Support simple and effective selection
- **R7 Provenance** – Store favorite options for future comparison
- **R8 Comparison** – Support criterion-wise comparison of options
- **R9 Details** – Enable direct reading of raw criterion values
- **R10 Sensitivity** – Show design parameter values of options
- **R11 Transparency** – Support awareness of intermediate decisions
- **R12 Accessibility** – Make the choice accessible to customers
- **R13 Export** – Provide the data of the chosen design for production

6.2 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 6.1. A mechatronics scientist at Linz Center of Mechatronics[1], who had about ten years of 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 [185]. 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 6.1a. At its core, it contained an augmented scatterplot, where points were interactively shown as radial bar charts [49] 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 6.1b)

1 www.lcm.at

and iteratively developed it towards a parallel coordinates visualization, whose interactions better reflect the optimization operations performed by the engineer (Figure 6.1c). 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 6.1d).

We performed a downstream validation against the threats of problem mischaracterization and wrong abstractions [185]. 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. The evaluation details are discussed in Section 6.4. We also validated the domain characterization by following up on the experts adoption of the tool for their daily work. The details of the adoption study are discussed in Chapter 8.

6.3 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. PAVED is publicly available at `paved.iva.igd.fraunhofer.de`. While it has been designed for engineers, our approach generalizes to any multi-attribute choice 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 decision 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 drill-down. 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).

6.3.1 *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 options (**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 decision 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 judgment of the decision-maker. Each criterion is thus meaningful for the evaluation of options 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, Oral et al. found that a commonality among many decision-making applications is the need for transparency in the sense of being able to directly retrieve decision criteria and options [195]. To make a multi-attribute choice, users need to be able to visually retrieve any criteria value from any option 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 options and criteria without the need for aggregation or selection of a data subset to view.

6.3.2 *Parallel Coordinates View*

PAVED's primary view shows a parallel coordinates visualization (Figure 6.4). 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

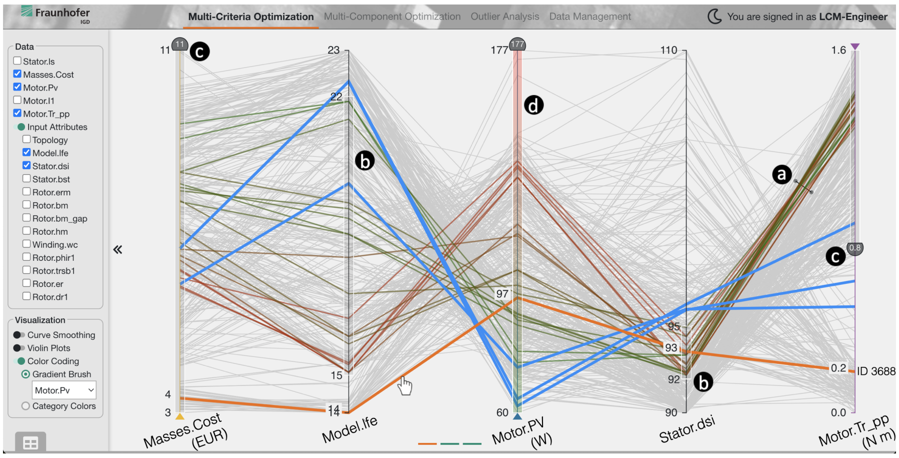

Figure 6.4: This image shows different interaction modes that drive the decision process. Standard brushes include a line brush for open exploration (a) and range brushes at the parameter axes (b). For criterion axes, we propose preference brushes (c), which are locked at the desired end of the axis 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 options (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).

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 [110] as our final visual encoding.

Parallel coordinates present a compact and lossless two-dimensional visual representation for multi-dimensional options. Different axis layouts have been proposed, e.g., many-to-many, force-directed, and three-dimensional layouts [117]. 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 [10]. Instead, we mark the desired direction of change by a triangular indicator at the respective end of the axis [198].

Motivated by the need to scale with hundreds of depicted options, our focus was on enhancing the perception of trade-offs and individual options. 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 options 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 [83]. Finally, each motor design is associated with a topology as categorical metadata. An effective technique to support the perception of nominal data is *color-coding* [96]. Categorical dimensions can also be used 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 attribute 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 after releasing an axis [96]. 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.

6.3.3 *Interaction Patterns*

Interaction is essential for an effective use of parallel coordinates. Selecting a subset of favorite options for detailed analysis (**R$_4$** Trade-offs) is complemented with filtering options according to performance constraints and preferences (**R$_5$** Filter). From a technical point of view, this corresponds to a selection by items versus a selection by attribute values. Available interactions are indicated by a transformation of the mouse cursor.

Selecting Favorite Alternatives

To support the user in scanning through the options, we provide *hovering* as the most basic interaction (Figure 6.4, orange). From the label to the right end of the hovered polyline, users can retrieve the option ID. This allows them to join insights about the alternative in focus with, e.g., offline data. To select a group of options, we provide a *line brush* (Figure 6.4a), which makes it easy to isolate options with particular characteristics (**R$_6$** User Interaction).

Users can *flag* alternatives by clicking on a polyline. Flagged options are permanently visible, even if they are not part of any other selection (Figure 6.4, blue). This enables a direct comparison with respect to each of their attribute values (**R$_8$** Comparison). They can also be

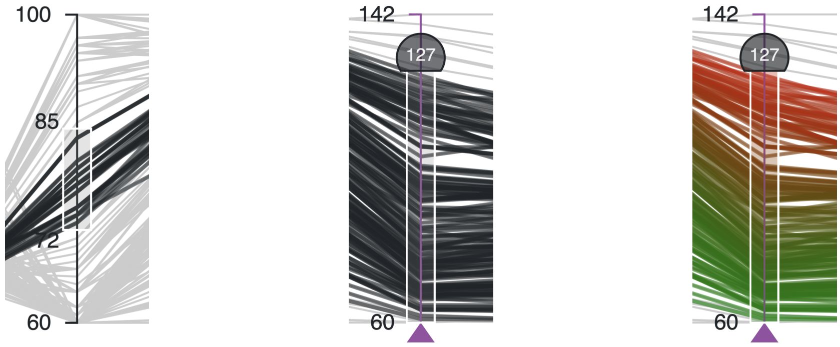

(a) Standard Range Brush

(b) Preference Brush

(c) Gradient Color Brush

Figure 6.5: 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).

considered the current result of the exploration. While exploring the Pareto front, the engineers flagged options 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 options in the form of range brushes applied to criteria and design parameters axes (Figure 6.5a). 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 6.4b) .

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 options 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 6.5b). Only the low-quality end of the brush can be dragged for filtering (Figure 6.4c). 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 enhanced by applying a red-green color gradient to its range that triggers a corresponding color-coding of the brushed polylines (Figure 6.5c). The *gradient color brush* was first introduced by Matković et al. [163]. 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 (**R3** 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 6.4d).

6.3.4 *Implementation*

The prototype presented in Figure 6.4 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* [36]. The web application, including the data management, runs on the client-side. The tabular data containing options and attributes is represented as an array of objects, with optimization direction and unit as optional meta data of attributes. Data can either 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.

6.4 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 [271]. 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.

6.4.1 *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 male experts other than our primary domain expert participated in the study: the major contributor of the tool currently used by the engineers in his role as a team manager as well as one simulation expert and three experienced motor designers in their roles of research and development engineers. The participants were aged between 30 and 50 with a working experience in their current or a similar role of five to eleven years. They reported to use visualizations quite often for exploring drive designs (average: 4/5, range: 3-5), with an average experience in developing interactive visualizations of 3.6/5 (range: 3-4) and an average familiarity with parallel coordinates of 3/5 (range: 2-4). A rapport was established during a preceding introduction in information visualization where all participants were present.

The notes taken during the think-aloud sessions were analyzed using a qualitative coding methodology [260]. 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 [138]. 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 [232]. 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.

6.4.2 *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 abstrac-

	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 6.1: Results of the System Usability Scale [232] 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.

tions in the domain characterization stage. We provide more detailed reflections in Sections 6.5.1 and 6.5.2.

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).

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. 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 [18]. 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 [232]. We present the individual scores broken down by question in Table 6.1. 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.

6.5 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 [236]. 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 [168]. 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.

6.5.1 *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. [13], 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 criterion value of any design from the visual encoding (**R_2** Overview) and 2) the transparency of the decision process (**R_{11}** Transparency). The latter is important because a comprehensible decision process *"helps to prove plausibility and justify the decision"* (A3). This is related to the need to communicate and collaboratively revisit a decision. Engineers need to explain how they arrived at 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 options 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 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 [69]. We were also pointed towards optimization scenarios involving interdependent components. We address this type of decision problem in Chapter 7.

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 optimization problems from different domains. 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.

6.5.2 *Visualization Idioms*

Motivated by the recommendation that *"studies [of parallel coordinates] in new application areas should be encouraged"* [117], we discuss aspects of our visual encoding, interaction techniques, and envisioned analysis workflow 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"* [117]. 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-attribute options. They particularly appreciated the brushing mechanism and observing its direct effect on the selection of options.

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. Still, it might not be effective for categorical decision criteria in its current form.

The existing analysis tool for the design and optimization of electric motors is built around an interactive scatterplot matrix depicting pairwise trade-offs. Scatterplots are known to convey correlations more effectively than parallel coordinates [146]. Our experts also stated that they prefer to observe pairwise trade-offs in a scatterplot, because it allows them to *"observe how much the Pareto front is bent"* (A5). 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 [305]. In their study, Dimara et al. found that tabular layouts were preferred over parallel coordinates for decision tasks [64]. However, tabular layouts often require users to explicitly express their preferences for ranking purposes [85]. Still, some experts confirmed the relevance of tabular visualizations: they would appreciate 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 options. 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.

Further investigation is needed to prove the applicability of PAVED in a wider context. A comparison of PAVED to other visualization techniques on different data sets, e.g., regular parallel coordinates, would add to the evidence of its impact beyond the individual use case. Our motivation behind not including this comparison was that the domain usefulness of PAVED needed to be validated before investigating its generalizability to other applications.

6.5.3 *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 investigate this problem in Chapter 7.

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.

6.6 CONCLUSION

This chapter presented a design study on Pareto front visualization supporting engineers who are tasked with the parametric design of electric motors. Visualization research often targets data analysts, i.e., people who are interested in understanding the data, not make decisions with data [68]. To the best of our knowledge, this design study was the first to explicitly address decision-makers and multi-attribute choice over analytic tasks at the time of implementation.

In close collaboration with the domain experts, we studied their needs, tasks, and goals related to choosing the most preferred motor design. The identified requirements guided our iterative development of PAVED, an interactive parallel coordinates visualization for exploration of design options. The multi-attribute options 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 increases transparency of the trade-offs involved, which is essential, not only for justifying the final choice to their customers (Section 2.2.3).

A qualitative field study suggests the effectiveness of our visualization as part of a real-world engineering design workflow. This is also reflected in the results of a quantitative usability testing, which resulted in a usability well above average. This outcome adds evidence to the benefit of introducing a well-known visualization like parallel coordinates in industry applications [117]. Finally, we reflect on the domain-specific problem characteristics, visualization design, and methodological considerations, which adds to the body of real-world experience with studying multi-attribute choices. In realizing that we had missed a shortcut, we particularly learned to value the domain experts' understanding of their own visualization needs. This, again, confirms the need to study visualization support in real-world settings with users who are, in fact, decision-makers [68].

Although designed for engineers, our visualization generalizes to any multi-attribute choice where the preferences cannot be quantified in advance. Thus, it can also be used by other decision-makers like consumers or professional buyers, policy makers, or event managers. Different perspectives on the decision space could help align with a decision-maker's mental model, for example extending the abstract parallel coordinates view with 3D model views of the subject to decide upon [195]. Chapter 7 will investigate how the parallel coordinates visualization can be further adapted to also support decisions where multiple choices affect each other.

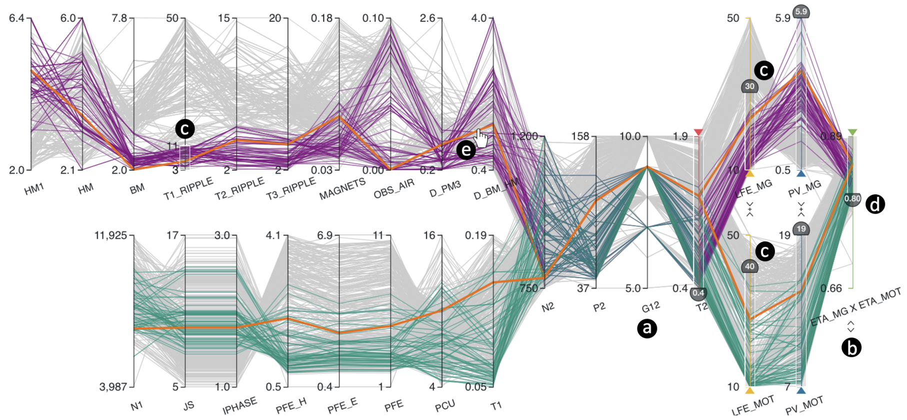

Figure 7.1: Composite Parallel Coordinates help users make decisions about multi-component systems, e.g., choose one of the 392 combinations of gear (top) and motor (bottom) alternatives for an electric drive. Axes can be merged to represent shared properties like the gear ratio (a) or to compute system performance indicators from component criteria (b). Filters can be applied to single components (c) or to the entire system (d). Hovering a gear alternative highlights it together with the compatible motor alternative (e).

7

COMPO*SED: COMPOSITE PARALLEL COORDINATES FOR CO-DEPENDENT MULTI-ATTRIBUTE CHOICES

THIS chapter proposes a novel parallel coordinates technique to effectively represent the interplay of two or more multi-attribute choices that affect each other. For a single multi-attribute choice, interactive visualizations have been proposed around helping decision-makers deal with sets of alternatives, conflicting attributes, and informal subjective preferences [175]. With the parallel coordinates plot [109] being an important representative, Chapter 6 presented a design study on its use in engineering design as an important application area. While the design study demonstrated the usefulness of parallel coordinates for motor design, its evaluation unveiled that what the domain experts typically need to decide upon are not single but multiple interacting mechatronic components (Section 6.5.1). This task had also been identified in a previous design study on car engine design [28]. Where multiple components are operated together, decision-makers face the task to work through a series of choices. They can affect each other in two ways. First, making one choice might rule out options of another choice (e.g., deciding for a camera body rules out lenses that have a different mount type). Second, the overall performance can only be judged with all intermediate choices together, because it is an emergent property of their composition and cannot simply be computed from individual performances. The challenge for visualization design is to allow for trade-offs regarding both the individual choice goals as well as those of the overall decision and to allow decision-makers to flexibly move between these two perspectives. This is complicated by the attribute sets, on which alternatives of the different choices are defined, not necessarily being comparable (see Section 4.3). Grounded in but not limited to a continued collaboration with the domain experts from the design study in Chapter 6, this chapter therefore turns towards the problem of making co-dependent multi-attribute choices.

Section 2.4.2 concluded that systems engineering design requires a simultaneous exploration across individually modeled but dependent components. This chapter investigates how parallel coordinates can be adapted to move from single multi-attribute choices to co-dependent choices. As co-dependent choices have not been addressed by visualization researchers before, we develop a characterization of the targeted problem. From there, we investigate layout strategies to depict both individual options at the component level as well as the dependencies between options of different components at the system level. The investigation also includes dedicated interaction patterns to strengthen the perception of component relationships and support efficient navigation through the large solution space. As a result, we present Composite Parallel Coordinates to help decision-makers choose the most preferred design alternative of a multi-component system. Unlike commonly employed multi-component models or iterative trial-and-error processes, it supports a simultaneous exploration of the components involved. The visualization technique presented in this chapter has been previously published [52].

What the reader can expect from this chapter:

- A characterization and abstraction of a new type of decision task called co-dependent multi-attribute choices (Sec. 7.1).
- A review of visualization works on analyzing related data sets, suggesting that a simultaneous analysis of individual items and their conditional combinations is not yet supported (Sec. 7.2).
- The design of *COMPO*SED*, a parallel coordinates technique to explore co-dependent choices for the first time (Sec. 7.3).
- Two usage scenarios and a case study that prove the applicability of COMPO*SED for investigating the side effects between intermediate decisions when making trade-offs (Sec. 7.4).
- An observational study based on a natural conversation of two experts jointly using COMPO*SED suggesting that it is suitable as an alternative to artificial think-aloud monologues (Sec. 7.4).

7.1 PROBLEM CHARACTERIZATION AND ABSTRACTION

Section 2.4.2 described the task of devising a system that meets desired properties within given constraints. The essential – and challenging – characteristics of a system are its combinatorial nature and the interactions between its components that lead to emergent properties. Emergent properties are properties of the system that the individual components do not possess when acting separately [249]. The characteristics pose three major challenges for systems design:

CH_1 **Combinatorial optimization** typically entails a huge solution space, even if restrictions apply. This prohibits an assessment of all possible system designs.

CH_2 **Interoperability constraints** restrict how components can be connected in a system. Consequently, individually optimal component alternatives might not be interoperable.

CH_3 **Emergent properties** make the system performance difficult to derive from individual component performances. In particular, local optimality might not yield a global system optimum.

This section develops a detailed characterization of the targeted problem. We abstract the interaction of subjects to decide upon in the form of a data model (Section 7.1.1). From the common difficulties and needs associated with system design problems, we then extract analysis tasks that make up the decision process (Section 7.1.2).

7.1.1 *Data*

The subject under investigation is a fixed set $C = \{C_1, ..., C_k\}$ of components that together form a system. Each component is optimized individually based on multi-run simulation, which describes the performance of the component under different design parameter settings. This results in one simulation ensemble per component. Each member of a component ensemble V is an alternative of this component. The previous design study provided an abstraction for alternatives as resulting from a multi-run simulation model (Section 6.1.2). As a reminder, a *component alternative* can be formally described as the union $\mathbf{v} = (\mathbf{x}, \mathbf{y}) \in V$ of a design option $\mathbf{x} \in X$ and its simulated performance $\mathbf{y} \in Y$ with design parameters X and criteria Y.

Although the components seem independent at first, they need to be integrated to achieve the purpose of the system. Thus, the component ensembles cannot be analyzed separately. Instead, an optimal combination of component alternatives requires a consideration of two levels: the component and the system level (Figure 7.2).

At the system level, the alternatives available for each component are put into the system context to account for their *interoperability* (**CH_2**), i.e., components can only be combined under certain conditions, and

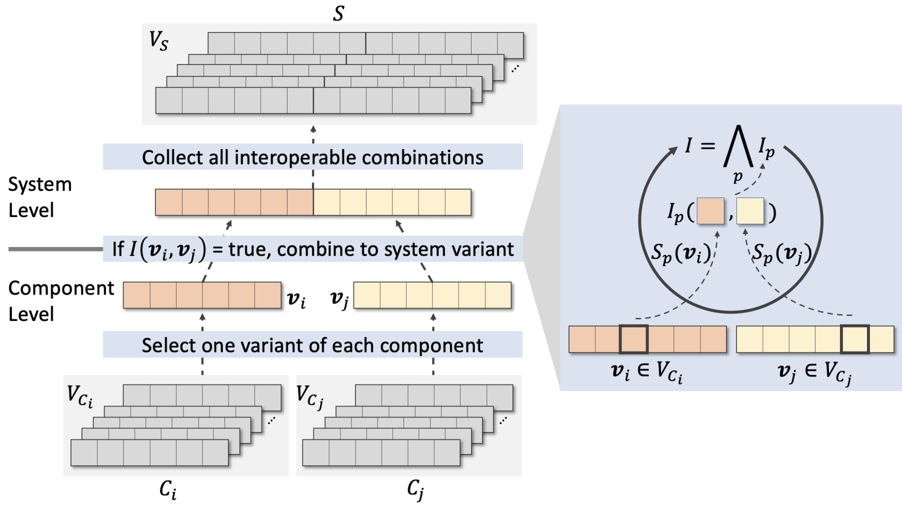

Figure 7.2: The system-oriented data model, going from the component level (bottom) to the system level (top). Only interoperable component alternatives (right) can be combined. The resulting system alternatives are augmented with global criteria approximating their emergent performance.

emergent properties (**CH3**), i.e., system performance is a synergy of component performances. The following formal description reflects these two dependencies in the data model.

Interoperability

To form a smoothly operating system, each individual component needs to fit its neighbors, physically and functionally. Mittal and Frayman use the idea of ports to describe such intercomponent boundaries [179]. A port is where a component connects to other components. Since we do not assume arbitrary connectivity, a port is also associated with constraints. For example, a lens can only be mounted on a camera body with a fitting mount type.

To represent ports and their constraints in our data model, we looked at how data sets are joined in relational databases [56]. A join condition specifies whether items from different data sets can be combined into a single type. In our case, items are alternatives of different components that are combined into a system alternative based on an interoperability condition. A system alternative is valid if it contains exactly one alternative of each component in the configuration. To avoid potentially incomplete system alternatives, we use inner joins to model the interoperability of components. Inner joins consider a tuple of component alternatives as a system alternative, if and only if all alternatives match the given condition.

Consider two components C_i and C_j with alternative ensembles V_{C_i} and V_{C_j}. The inner join computes an ensemble V_S of system alternatives, i.e., an ensemble containing combinations of interoperable alternatives from C_i and C_j (Figure 7.2):

$$V_S = \{(\mathbf{v}_i, \mathbf{v}_j) \in V_{C_i} \times V_{C_j} \mid I(\mathbf{v}_i, \mathbf{v}_j) = \mathit{true}\} \tag{7.1}$$

The interoperability condition is represented by a generic predicate function I. As interoperability is concerned with design space restrictions, I is evaluated on the components' design parameters. We detail the definition of I based on the following assumptions:

- Interoperability might be constrained by more than one port. We distinguish the involved ports by an index p. All defined port constraints must be met for I to yield *true*.
- Each port constraint is described by a boolean predicate I_p that is defined on two design parameters: one parameter of C_i and one parameter of C_j. We write $S_p(\mathbf{v}_i)$ and $S_p(\mathbf{v}_j)$ to select the values of the two design parameters from the respective component alternatives $\mathbf{v}_i$ and $\mathbf{v}_j$.
- The predicate I_p of each port is selected from a class of boolean functions during the specification of the system configuration.

Given the first two assumptions, we define the predicate function I as the logical AND operation ($\bigwedge$) of all individual port predicates I_p:

$$I(\mathbf{v}_i, \mathbf{v}_j) = \bigwedge_p I_p\left(S_p(\mathbf{v}_i), S_p(\mathbf{v}_j)\right) \tag{7.2}$$

By specifying the predicates I_p and the design parameters they operate on, I can be chosen to account for a variety of interoperability conditions in various domains. The possible functions for I_p can take two different forms: the two design parameter values are either directly compared or they are aggregated and compared to a constant c.

$$I_p(x, y) = \begin{cases} x \mathbin{◐} y, & \text{(direct comparison)} \\ x \odot y \mathbin{◐} c, & \text{(comparison with constant)} \end{cases} \tag{7.3}$$

Above, $◐ \in \{=, \neq, >, \geq, <, \leq\}$ and $\odot \in \{+, -, *, /\}$. Thus, Equation 7.3 describes 30 functions that can be used to define the most common port constraints during the system configuration. A natural join or a theta-join [56] is implemented through a direct comparison (natural join) or a comparison with a constant (theta-join) on the relevant design parameters. For example, for the natural join of camera bodies and lenses, we can choose I_p to operate on the components' design parameters *mount type* using "=" as the comparison operator ◐. Depending on the interoperability conditions for a system design problem, the definition can be extended with custom predicates if they are using exactly one design parameter of each component. In this way, our technique generalizes to a wide variety of applications.

Emergent Properties

This dependency between components refers to the system performance as an emergent property of a combination of component alternatives. We distinguish between local criteria capturing the individual component performances and global criteria capturing the emergent system performance of valid system alternatives. Following the formal characterization of Weidele [293], this can be described as conditional data: if component alternatives meet the interoperability predicate, the resulting system alternative can be augmented with details about the system performance. Instead of costly simulations, we approximate the system performance by computing global criteria from selected local criteria. A computation based on semantically related local criteria may involve simple mathematical operations, like adding up individual component costs to a system price but may also use more complex functions to address non-trivial compositions. However, the global criteria can also be computed by a conventional weighting approach to aggregate local criteria without any semantic relationship.

7.1.2 *Tasks*

The design of engineered systems is relevant in a variety of domains, each providing different environments regarding specifications, domain knowledge, and existing workflows. Wang and colleagues present six analytic tasks related to inferring meaningful information from ensemble data [285]. Their compilation does not cover the main task of system design, which lies in an informed *combination* of members of different ensembles. The primary goal of the analyst is to determine the most preferred combination of component alternatives, such that the resulting system's performance is optimized with respect to a number of attributes. As there is no inverted simulation model that tells analysts how to choose the design parameters to achieve the desired performance, the available combinations need to be explored [255]. The exploratory choice involves multiple co-dependent multi-attribute choices: for each component involved in the system configuration, analysts need to select the best option among a finite number of alternatives, such that their interplay is optimized. Herein lies the main challenge of system design. In fact, a central aspect of system design is the "*subordination of individual goals and attributes in favor of those of the overall system*" [137]. Such mutual dependencies significantly intensify the decision-making process. Deciding for the best combination turns into a series of choices, whose potential side effects need to be considered when making trade-offs. An exploration across ensembles is thus guided by the following questions:

- **T1 Overview** – What are the value distributions of attributes at the component level and at the system level?

- **T_2 Competition** – What is the nature of conflicts among criteria? How important are component performances and system performance?
- **T_3 Filter** – What is the effect of system and component constraints?
- **T_4 Subjective Evaluation** – Is an alternative feasible? Does it balance the criteria according to the stakeholder's preferences, tolerances, domain knowledge, and experience? Which alternative is superior?
- **T_5 Connectivity** – Given a (set of) component alternative(s) of interest, which alternatives of the remaining components are interoperable?
- **T_6 Navigation** – How does deciding for a component alternative affect the availability of remaining alternatives and the system performance in reach? Which alternative improves the emergent system properties?
- **T_7 Key Component Alternatives** – Which component alternatives yield a good performance, while being interoperable with many other alternatives?
- **T_8 Partial Choice** – Does a partial component choice, i.e., selecting a set of possible component alternatives, leave enough room for balancing the properties of the remaining components?
- **T_9 Alternative Replacement** – How does a component alternative contribute to the system performance? What is the effect of replacing the component alternative? Which replacement increases component or system performance?

7.2 VISUALIZATION OF MULTI-MODEL DATA

For depicting multi-attribute variants of a single component, we can draw from existing works on multivariate Pareto front visualization (Section 3.2.1). Jointly analyzing variant ensembles of more than one component relates to visualization approaches addressing a simultaneous investigation of different but related data sets (Section 7.2.1). Since we use parallel coordinates, we also investigate related approaches for organizing parallel coordinates axes (Section 7.2.2).

7.2.1 *Visualization of Multiple Related Data Sets*

System design requires an investigation of related data sets representing components and their interactions. Konyha and colleagues conclude that single table approaches are insufficient to describe such data and its dependencies [133]. This view is shared by Kehrer and Hauser, who identify multi-model scenarios resulting in two or more interacting data parts as a promising direction of visualization research [122]. A central question in such cases is how to investigate patterns across data sets.

Coordinated multiple views may link multiple tables via primary and foreign keys as done, e.g., in *Snap-Together* [190]. Liu et al. consider the relationships between data items as a graph and propose the system *Ploceus* for a network-based visual analysis [150]. With *Domino*, Gratzl et al. propose a meta-visualization technique allowing users to create explicitly linked views to represent data subsets and four degrees of relationships between them [84]. Working with multi-resolution models, Spletchna et al. address the complication of only partially overlapping parameter spaces as a key challenge [259].

The links between our component data sets are not defined by shared identifiers but by value predicates. Kehrer et al. propose a similar abstraction of the relation between two data sets, which they call *interface* [123]. Their abstraction addresses multi-model data in a spatial domain, which can be exploited to describe the relations via co-location. Closest to our approach is the work by Splechtna et al., who propose to relate items of different data sets based on their properties regarding one or multiple (common) attributes [258]. While we build upon such conditions to represent the interoperability of components in a system, their approach cannot handle emergent properties.

The analysis of emergent properties plays a significant role for the optimization of complex engineered systems. Basole et al. propose a network-based visual analytics tool for system design that explicitly considers how intermediate decisions influence system-level properties [22]. While their approach focuses on an iterative reconfiguration of the system, our approach is based on systematic sampling of the design space upfront to gain a broad overview early in the process. Closest to our work, Marth et al. use scatter plots to evaluate the joint performance of a motor and a gear in a side-by-side arrangement [161]. They provide performance criteria for individual motor and gear variants as well as for their combinations by summing up the individual criteria (e.g., the sum of losses or lengths). We generalize the specification of interoperability and joint performance and visualize these system-level properties together with component-level properties. So far, no multi-ensemble approach has been proposed that allows for a simultaneous exploration and optimization of individual ensemble members as well as conditional combinations of ensemble members.

7.2.2 *Axis Configurations for Parallel Coordinates*

In their survey, Heinrich and Weiskopf define a composite parallel coordinates plot as a composition of several visualization layers, e.g., axes and brushes [96]. In contrast, we propose a side-by-side composition of multiple parallel coordinates plots, which emphasizes the dependencies and emergent properties in multi-component systems.

The core challenge of our parallel coordinates composition is the arrangement of axes. In the conventional layout, every dimension has two direct neighbors. A strategy to overcome this limitation for visualizing many pair-wise relations is to replicate axes of individual dimensions. Lind et al. combine multiple axis orderings in a many-to-many parallel coordinates plot [147]. Replicated axes depicting the same dimension are arranged in polygons to communicate all pairwise relations of dimensions in a non-overlapping way. Claessen and van Wijk propose a generalization where axes can be positioned freely and linked via scatter plots (orthogonal axes) or parallel coordinates (parallel axes) [55]. We did not consider such an approach because it results in a complex visualization layout, even for single components.

We want particular relations across components to stand out. Multiple strategies have been proposed to visually aggregate dimensions with similar semantic meanings. Andrews et al. introduce aggregate axes that replace related dimensions by substituting the dimension values with their mean [9]. Axes can be interactively collapsed and expanded. Bhattarai and colleagues use the sum to merge dimension axes for an exploration of material compositions [33]. Garrison et al. aggregate dimensions that contribute similarly to the variance of a data set by mapping parallel coordinates axes to the first and second principal components of the dimensions [78]. In their product comparison tool *ConfigurationFinder*, Riehmann and colleagues organize semantically related dimensions in groups that are represented by an expandable proxy axis [222]. The approaches mainly differ in how the related dimensions are identified (domain knowledge or automated analysis) and how the aggregations are calculated and presented. In our case, related dimensions of different components are derived from the analyst's knowledge. Their axes can be merged using different functions to depict system-level properties. A primary challenge with our approach is to represent the interoperability conditions.

Conditional parallel coordinates by Weidele use predicates to insert nested axes for conditional dimensions that apply only to data sets satisfying specific properties [293]. While this approach allows for representing items of different types in one view (e.g., cameras and lenses), it does not enable a combination of items of two or more types. In our approach, all dimensions can be shown from the beginning because the predicates imply combinatorial constraints, leaving the dimension schema unchanged. In particular, axes unique to one type (i.e., component) are shown at the same level of detail as shared axes.

Regarding the ambition to visualize multiple ensembles, most similar to our approach are the nested parallel coordinates proposed by Wang et al. [286]. They use nested axes to compare data distributions from multiple ensembles that originate from climate simulation at different resolutions. Our Composite Parallel Coordinates are inspired by their approach, i.e., to assemble juxtaposition and superimposition for analysis within and between different ensembles. However, while their ensembles provide different resolutions of the same subject, the components represented by our ensembles are different subjects.

7.3 DESIGN OF COMPOSITE PARALLEL COORDINATES

Based on our preceding characterization of co-dependent choice data and tasks, we propose COMPO*SED, a novel variant of parallel coordinates for their visualization. COMPO*SED allows analysts to explore the possible combinations of component alternatives while taking into account both the individual component properties as well as their emergent system properties. Parallel coordinates offer a compact and lossless two-dimensional visual representation for multi-dimensional observations. We made them the basis of our visualization design primarily for their lossless mapping and flexible axis arrangements [55] but also for their simple applicability and wide-spread use in multivariate data exploration. From the results of our previous design study (Chapter 6), we were confident that parallel coordinates are accessible for analysts performing single multi-attribute choices.

System design requires an observation of design options at the component level and the system level. This leads to three conflicting design goals (**G_1**, **G_2**, **G_3**) and one independent goal (**G_4**):

- **G_1 Component-level analysis** requires a stand-alone observation of individual properties per component. Without prior relevance information, all components and properties are considered equally meaningful for analysis.
- **G_2 Context awareness** requires to relate observed component properties to semantically similar properties of other components.
- **G_3 System-level analysis** needs an observation of similar properties across components. The evaluation of interoperability and emergent performance benefits from explicit system properties.
- **G_4 Layout stability** is an overarching design goal. In contrast to open data exploration scenarios, where no analysis strategy is imposed, system design relies on a clear mental model of the system structure and properties to investigate. A stable overview layout allows analysts to focus on trade-offs instead of adapting to varying positions of components and their properties.

Conventional parallel coordinates lack the ability to depict the dependencies between individual ensemble members (Section 7.3.1). We therefore propose Composite Parallel Coordinates (Section 7.3.2), whose axis layout (Sections 7.3.3 and 7.3.4) and interaction patterns (Section 7.3.5) reflect the notion of a system being a composition of interacting components.

7.3.1 *Reviewing Conventional Parallel Coordinates*

A conventional parallel coordinates plot depicts a single multi-dimensional data set, where all items are defined in the same variable space. To visualize multiple component ensembles, their variable spaces

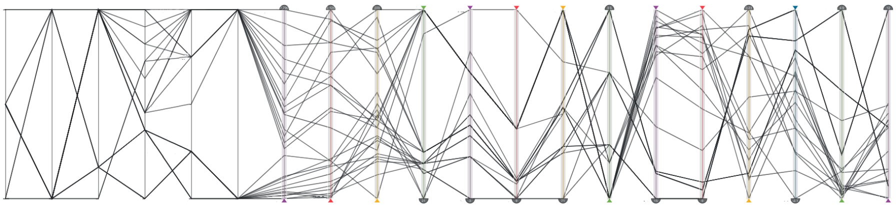

(a) Conventional Parallel Coordinates

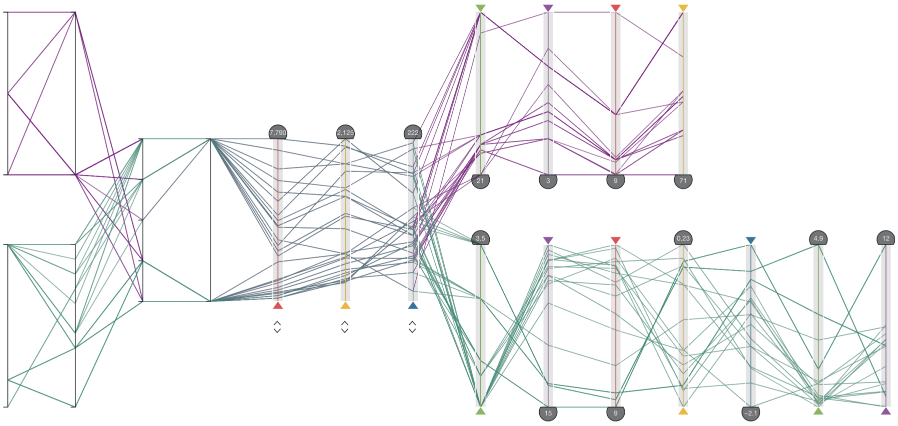

(b) Composite Parallel Coordinates

Figure 7.3: Two-component systems depicted by (a) standard parallel coordinates (PC) vs. (b) the proposed Composite Parallel Coordinates (ComPC). PC display the combinations over their joint variable space, while ComPC partially juxtapose the components. In contrast to PC, they explicitly depict the combinatorial problem.

need to be merged during the data transformation step. This can be achieved by joining the component alternatives according to their interoperability. Due to the combinatorial nature, the number of items and variables to visualize increases drastically (upper bound $n * m$ for items or $n + m$ for variables). Conventional parallel coordinates then result in a plot with many side-by-side axes, where the polylines represent the ensemble V_S of system alternatives (Figure 7.3a).

In this plot, the subdivision of system alternatives into individual component alternatives is not obvious. This makes it difficult to perceive how the properties of the system originate from the interactions between the individual components. It complicates the central aspect of system design, i.e., the subordination of component-specific characteristics in favor of the system performance [137]. The root of this problem is the axis layout being restricted to the horizontal direction.

On the one hand, the success of a choice at the component level is determined by the component properties (**G1**). The evaluation is facilitated if all axes of the same component C_i are placed directly next to each other, such that each component alternative $\mathbf{v}$ is represented as a self-contained polyline. This axis order supports tasks like gaining an overview of component alternatives (**T1**), determining key component alternatives (**T7**), or replacing a component alternative (**T9**).

On the other hand, determining the success of an intermediate choice builds upon an evaluation of emergent system properties (**G2** and **G3**). This can only be achieved by placing the axes belonging to different components C_i and C_j directly next to each other. Such an axis order supports tasks like determining interoperability (**T5**), evaluating the system performance (**T4**), or navigating the combinatorial design space (**T6**). However, it contradicts the component-wise adjacency of axes required for evaluating the choice at the component level.

Using conventional parallel coordinates, **G1**, **G2** and **G3** can only be achieved if we allow the axis order to be interactively adjusted to the varying analysis focus. However, this would mean violating the requirement for layout stability (**G4**).

To summarize, conventional parallel coordinates are not suited to meaningfully represent both component membership (**CH1**) as well as interoperability (**CH2**) and system performance (**CH3**). This boils down to them being restricted to depicting a single ensemble of system alternatives using a one-directional axis layout. Parallel coordinates cannot communicate the dual role of component variables that contribute to both component-specific properties as well as system-wide properties.

7.3.2 *General Visualization Design*

As explained previously, conventional parallel coordinates do not allow for an understanding of how the components work together. This issue can be solved by visualizing the components individually instead while paying particular attention to the interfaces between them. We map the component ensembles to distinct parallel coordinates plots. Depicting their interfaces poses an inherent challenge when visualizing a system as a composition of components. It requires careful integration of the different parallel coordinates plots into the same view. Javed and Elmqvist define this approach as *composite visualization* [115]. We make use of their design space to convey the idea of Composite Parallel Coordinates. Our visualization design was guided by the following question: how to make component properties (**G1**), their context (**G2**), and system properties (**G3**) equally accessible in a stable layout (**G4**)?

A composite visualization is a natural choice for depicting a system of interacting components. In contrast to the concept of coordinated multiple views [224], where different visualizations depict different aspects of the same data items, our composition involves multiple instances of the same visualization to depict different but related data items. The reason is that the primary task of making a multi-attribute choice is the same for each component in the system to be designed. Each system component ensemble V_{C_i} is depicted by one parallel coordinates plot.

How do these views become part of a composite visualization? Data-wise, they are independent because the variable spaces $X \cup Y$ are different for each component. The views' dependency originates from the domain-specific semantics regarding interoperability and emergent properties (see Section 7.1.1). The visual design task is to communicate these implicit dependencies as explicitly as possible.

The design space of composite visualizations proposes two symmetric and two asymmetric composition strategies [115]. The asymmetric strategies, namely overloading and nesting, impose an imbalance between views, which does not match the inherent symmetry of the system design problem, where all components are considered equally important. Thus, we turn toward the two symmetric strategies: juxtaposition and superimposition (Figure 7.3b). Juxtaposed parallel coordinates plots address the component level by depicting the largely different variable spaces. Superimposed axes of different plots address the system level by communicating emergent properties like interoperability conditions and overall system performance.

Our strategy accounts for the pairwise relations between the k variable spaces of the components. It involves different visual mappings to communicate the parts of a relation between variable spaces:

- **Shared** – Two variable spaces share parts where they exhibit common variables. Typically, these variables are design parameters considered for interoperability modeling, e.g., *gear ratio*.
- **Related** – The related parts of two variable spaces contain those non-common variables that contribute to interoperability and emergent properties. Related variables can be design parameters or criteria, such as *motor price* and *gear price*.
- **Unique** – Those parts of a variable space that neither involve common nor related variables are unique, e.g., *motor iron loss*.

Below, we describe our design choices regarding the visual mappings.

7.3.3 *Juxtaposition for Component Level*

At the component level, the decision-maker focuses on the individual properties of one component at a time ($\mathbf{G_1}$). Besides context awareness, considerations that involve other components, like interoperability and system performance, are of secondary importance. The simplest way of presenting an overview of all component ensembles is a juxtaposition of separate visualizations. Due to their independence, they allow analysts to focus on individual components without interference or distraction. As all components are considered to be equally important, we symmetrically divide the available visual space.

Juxtaposed views are generally highly flexible regarding their arrangement. However, in our case, the layout quality particularly depends on its ability to display semantically related properties of different components spatially close to each other in order to maintain

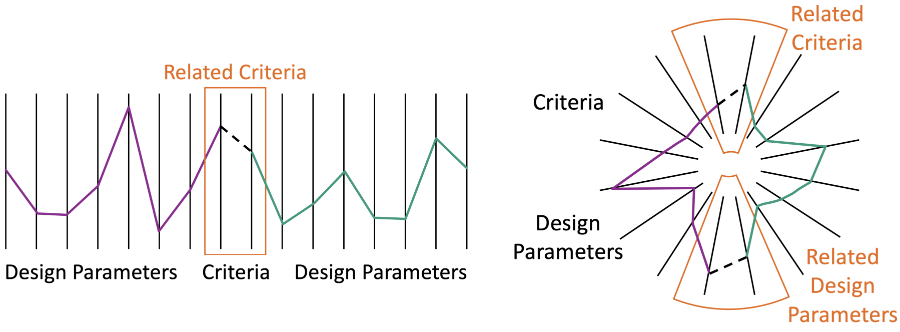

(a) Horizontal Layout

(b) Circular Layout

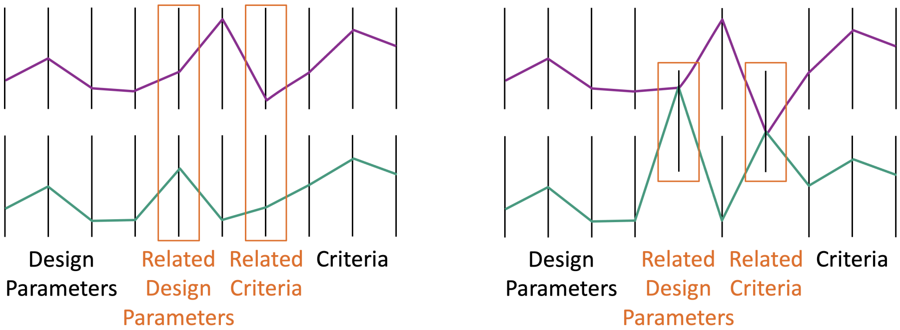

(c) Vertical Layout

(d) Vertical Layout + Superimposition

Figure 7.4: Different layout strategies to represent two components and their dependencies. Related properties of different components should be depicted close to each other. (a) Concatenated PCPs allow for adjacency of one pair of related properties. (b) Bending them to a circle adds another adjacent pair of related properties. (c) A vertical arrangement allows for more than two pairs of properties with similar semantics to be placed close to each other. (d) Where applicable, shared and related component properties can be merged to explicitly reflect system properties.

the system context (**G2**). The value of juxtaposition then stems from the boundary between two views conveying shared semantics. In the following, we discuss different layout options in light of this aspect.

Let us consider the composite visualization of two component ensembles V_{C_i} and V_{C_j}. A naïve approach would be to horizontally concatenate the two parallel coordinates plots (Figure 7.4a). Perceiving the system context requires the related properties of both components to be depicted close together. To achieve adjacency, the respective axes are placed at the inner ends of the plots. However, only a single pair of design parameters or criteria is adjacent, i.e., only one interoperability condition I_p or emergent property can be communicated.

The second option is a circular layout (Figure 7.4b). The related design parameters are placed at the inner boundary of the concatenated plots and the criteria contributing to emergent properties at their outer ends. The plots are then bent to a circle, such that the criteria, too, are adjacent. The result resembles a radar chart. Still, this layout conveys only one interoperability condition I_p together with one emergent property and cannot be extended to more than two components.

A third option is to arrange the parallel coordinates plots vertically (Figure 7.4c). The vertical distribution clearly separates the visual representations of individual components, thus enabling an efficient perception of the component level. The horizontal direction can then be exploited to position variables with similar semantic meanings but belonging to different components one below the other. In this way, the boundary between two views conveys shared semantics via multiple interoperability conditions and emergent properties and properly accounts for context awareness. In addition to that, the layout offers the potential to be extended to more than two components.

Based on the requirements imposed by the component-level analysis (**G1**) and context awareness (**G2**) together with the overarching layout stability (**G4**), the juxtaposition with vertical layout is the most promising option to proceed with. While the boundary between views conveys the system context via shared variable semantics, any linking between data items of different views is revealed only upon user interaction. Relations between alternatives $\mathbf{v}_i \in V_{C_i}$ and $\mathbf{v}_j \in V_{C_j}$ of different components are difficult to perceive. These relations describe the interoperability of component alternatives as well as their joint performance. They refer to a system-level analysis (**G3**), which is detailed in the following section.

7.3.4 *Superimposition for System Level*

Dependencies at the system level manifest in parts of the variable spaces being shared (common variables) or related (variables contributing to interoperability conditions or emergent properties). Superimposition means to overlay two plots in a single view [115]. We implement it by allowing the user to merge those axes that are associated with shared or related parts of the variable spaces (Figure 7.4d). As a result, the interface of two components is depicted by those polyline sections that intersect the superimposed axes and thus share the same visual space at the boundary between the juxtaposed plots. Unlike any other layout that we considered, this strategy solves the component-level analysis (**G1**) with context awareness (**G2**) and the system-level analysis (**G3**) while providing a stable layout (**G4**).

The design parameters and criteria that are unique to individual components are not affected by superimposition. They are represented by the juxtaposed atomic axes of the individual parallel coordinates plots because they do not semantically relate to another component's properties. An example is the property *focal length* of a camera lens, which remains unaffected no matter which camera body is chosen to be combined with it.

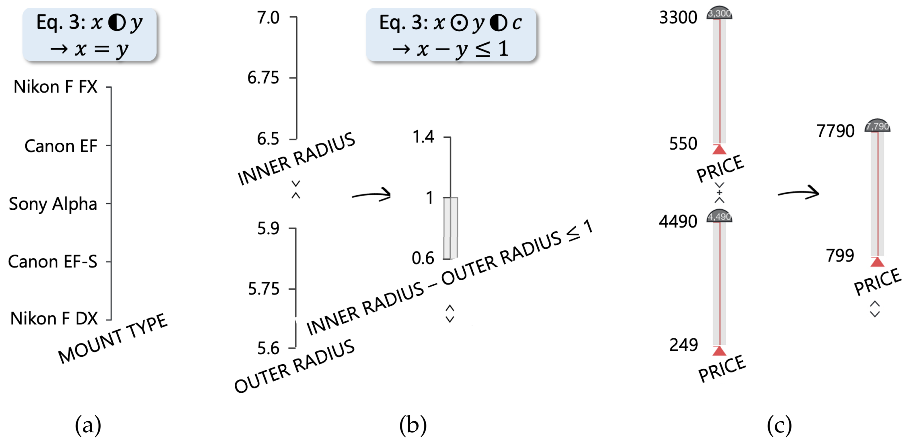

Figure 7.5: Composite axes represent interfaces between components. (a) Shared design parameters can contribute to interoperability (e.g., mount types should be equal). (b) Related design parameters specify interoperability via inequality (e.g., radii of components should differ by one unit at most). (c) Related criteria are added up to a system criterion.

For the shared and related parts of the variable spaces, we introduce *composite axes*. These axes can take different forms, depending on the represented type of dependency and the current analysis focus. What they have in common is that atomic axes of different plots might be collapsed into a derived axis. It is placed vertically in between the involved juxtaposed views and acts as an independent variable itself.

We distinguish three types of dependencies between components, which are represented by slight variations of composite axes:

- **Shared design parameters** are depicted using a permanently collapsed axis (Figure 7.5a). They are not expandable, as this would mean to duplicate the axis and thus add redundancy.
- **Related design parameters** are depicted using separate axes initially, but can be collapsed by applying a predicate function describing the interoperability condition I_p (Figure 7.5b).
- **Related criteria** are also depicted using separate component axes, but can be collapsed via a mathematical operation that maps the component criteria to a system criterion (Figure 7.5c).

Composite axes have a button located beneath them to collapse and expand the associated component variables. When two component axes are collapsed, the derived axis is inserted vertically centered between the two original axes, replacing them. The polylines of both involved plots are updated to intersect the collapsed axis. Where non-identical parameters are merged, this requires the creation of combinations. To keep complexity low, the range of the collapsed axis is computed naïvely from the extreme values of the original axes. When the collapsed axis is expanded again, the component axes are inserted at their original position in the plot, replacing the collapsed axis (**G4**). The polylines are updated again to intersect the separate component axes.

The dependencies, i.e., which parameters can be collapsed and how, are prescribed by the application domain. The axis pairs and collapse functions are specified a priori by users in the form of metadata of the data set to be analyzed. Up to now, the users have managed to do so without a dedicated user interface. Still, whenever needed, only development efforts would be required to provide a user interface to not only specify but also adjust the axis pairs and collapse functions.

Value Mapping

The value mapping of a composite axis depends on its type and collapse function. Composite axes displaying shared design parameters do not require a dedicated value mapping. The permanently collapsed axis simply displays the original parameter values of the alternatives across both components (Figure 7.5a).

In contrast to shared parameters, related parameters have the same semantic meaning but are not identical. They are initially observed individually using separate component axes. Based on their semantic relationship, these axes can be collapsed to reflect properties at the system level. This requires the combination of component alternatives using a dedicated value mapping that derives an aggregated system value from the two original component values. These aggregated values are displayed on the collapsed axis. To specify the exact mapping, we need to distinguish related design parameters and related criteria.

For related design parameters, the collapse function is taken from the pool of interoperability predicates (Section 7.1.1). As an example, the outer radius of a motor rotor is related to the inner radius of the surrounding stator according to an inequality constraint. A predicate I_p is applied in two steps. First, the aggregate of the two component values is computed using the $\odot$ operator. In case of a direct comparison, where no operator is involved, the predicate is rewritten to a comparison with constant 0, e.g., $x < y \rightarrow x - y < 0$. For the rotor and stator, $\odot$ might be defined as subtraction such that $r_{inner} - r_{outer}$ describes the clearance between both mechanical parts. This aggregate value is displayed on the collapsed axis. Its range is derived from the aggregate values across all combinations of rotor and stator alternatives. In a second step, the comparison operator $◐$ is applied as a filter. In this case, the clearance should not exceed the value one, so $◐$ is $\leq$ and any aggregate value less than or equal to one is brushed on the axis (Figure 7.5b).

As an example for related criteria, the individual prices of a camera body and a lens might be added up to reflect the system price. Criteria can only be collapsed if they are 1) both to be minimized or both to be maximized and 2) associated with the same or relatable units. Upon collapse, any meaningful mathematical operation might be applied to the original values of the two involved component criteria. The

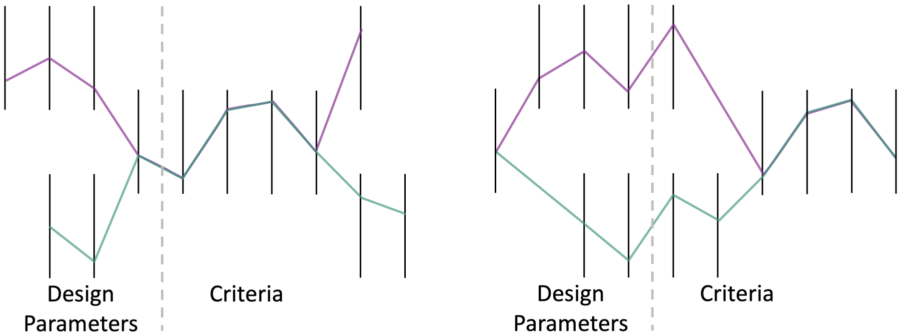

(a) Hourglass Shape (b) Wristwatch Shape

Figure 7.6: The axis order follows the input to output mapping. It minimizes alternation between separate and collapsed axes. Unlike the wristwatch shape (b), the hourglass shape explicitly depicts the bottleneck (a).

collapsed axis then depicts for example the total system costs as the sum of the two component prices (Figure 7.5c). Its range is computed by applying the same operation to the original minima and maxima of the component axes. This range covers all potential combinations of component alternatives but is not necessarily exploited.

Axis Order

The initial order of composite axes is generated by mapping the input-output order of the data model to the reading direction from left to right. The design parameters (input) are placed on the left side of the visualization, while the criteria (output) go on the right side. We can further divide the input and output into unique and shared or related variables. It should be noted, however, that strong alternation between separate and collapsable composite axes is not desired due to the turbulent polyline courses this generates.

Considering this, we can place the separate design parameters left, then the collapsable design parameters and collapsable criteria in the middle, and the separate criteria to the right (Figure 7.6a). The resulting shape resembles an hourglass. Alternatively, we can place the collapsable design parameters to the left, then the separate design parameters and separate criteria in the middle and the collapsable criteria to the right (Figure 7.6b). This shape looks more like a wristwatch. Within these constraints, the axes are initialized according to their order of occurrence in the data sets. We decide for the hourglass shape, because it explicitly communicates that the *bottleneck* of system design is the interaction between components.

Handling Line Overlaps

The fusion of component polylines at collapsed composite axes provokes line crossings that can make the course of individual lines hard to trace (Figure 7.7a). We developed two strategies to cope with this.

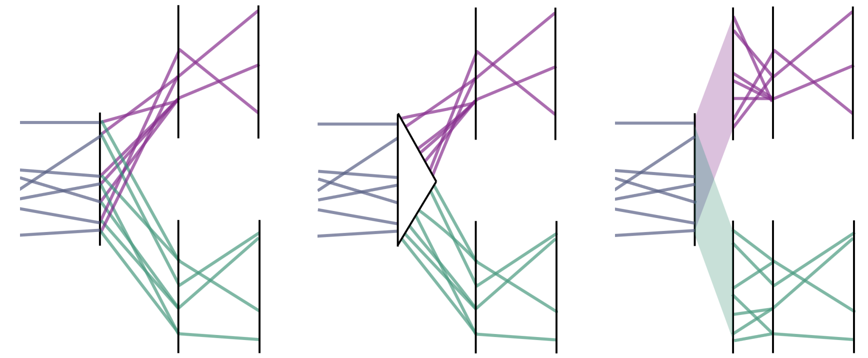

(a) Problem: Line Overlap (b) Solution: Triangle (c) Solution: Parallel Sets

Figure 7.7: Line crossings occur where separated polylines merge into one (a). Redirecting the polylines via duplicates of the collapsed axis using a triangle (b) or parallelograms (c) mitigates the overlaps.

The first strategy is inspired by the many-to-many parallel coordinates proposed by Lind et al. [147]. They arrange replicated axes of the same variable in a triangular shape to visualize a one-to-three relationship. We implement this concept to handle line crossings by replicating the collapsed axis, arranging the three identical axes in a triangle, and redirecting the polylines from the original collapsed axes via the two replicated axes to the adjacent separate axes (Figure 7.7b).

The second strategy resembles the parallel sets originally introduced for categorical data [136]. We use parallelograms to show the connection between the original and replicated axes (Figure 7.7c). The vertical orientation of the replicated axes matches well with the general layout. It explicitly communicates that the polylines split up into the two component levels. The values are easier to read from the axes, and the match between incoming polyline sections at the collapsed axis and outgoing sections at the replicated axes is easier to make. We thus opt for the parallel sets strategy to bypass the line crossings where separate and collapsable composite axes are adjacent.

7.3.5 *Interaction Patterns*

Interaction is essential for effective use of Composite Parallel Coordinates. It allows analysts to filter the available alternatives according to constraints and preferences as well as emphasize alternatives of interest. Our interaction mechanism involves three cascading selectors to gradually refine a selection of alternatives under focus.

The selectors are hierarchically structured: *filters* take precedence over *locks*, which in turn have priority over *mouseover* interactions. The outcome of any selector is a set of alternatives. Each selector operates on the outcome of the precedent selector: hovered alternatives are a subset of locked alternatives, which in turn are a subset of filtered alternatives. Note that this structure does not prescribe the order in

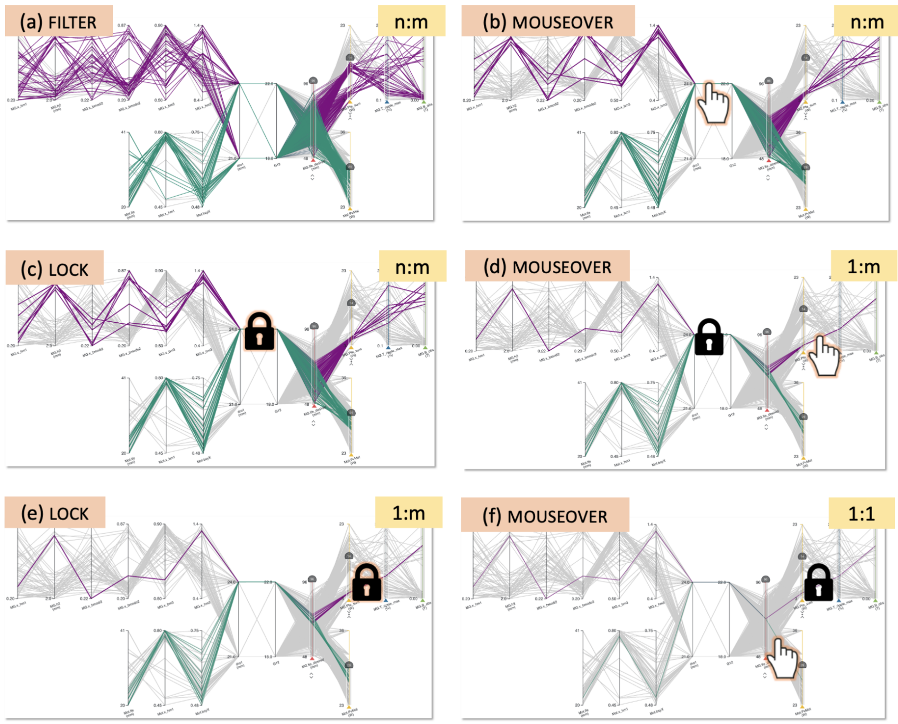

Figure 7.8: Selectors are applied at three precedence levels: filter, lock, and mouseover. The resulting selections range from many-to-many over one-to-many to one-to-one combinations. (a) A filter results in a set of alternatives that is refined by (b) a mouseover selecting system alternatives with a particular diameter and gear ratio. (c) Locking them allows for (d) mouseover exploration of the involved component alternatives. (e) One gear alternative is locked to explore the compatible motor alternatives. (f) One of them is finally chosen as the best fit.

which selectors have to be applied. Selectors are optional: if a selector is not active, the outcome of its precedent selector remains unchanged. At any time, the selection resulting from the cascade of currently active selectors is highlighted, while the remaining alternatives are depicted in grey for context. In the following, we describe the three selectors in the order of their precedence.

Eliminating Undesired Alternatives Via Filters

We provide filters in the form of range brushes that can be applied to any design parameter or criterion axis in the Composite Parallel Coordinates. As alternatives need to be evaluated regarding multiple constraints and preferences, multiple brushes can be combined into a composite brush using the logical AND operation. Where brushes represent interoperability conditions I_p, their composite brush corresponds to the overall interoperability I. Because it does not make sense to exclude alternatives with desired criterion values, brushes on criteria axes are tied to the high-quality end of an axis [54].

In large parts, filtering works in a standard way: alternatives covered by a brush are included in the selection. Eliminating undesired alternatives in this way results in a subset of acceptable options to proceed with (Figure 7.8a). However, the combinatorial nature of the optimization gives rise to some special considerations:

- A component alternative is selected either if it is brushed itself or if it can be combined with at least one brushed alternative of another component.
- A component alternative can be brushed itself on a unique axis or as part of a combination on a collapsed axis.
- When two axes are collapsed, the new brush position is determined by applying the collapse function to the original slider values. Upon expansion, the original slider values are restored.

Highlighting Desired Alternatives Via Locks and Mouseover

With a potentially large number of acceptable combinations remaining after filtering, users need support in scanning through the filtered alternatives to further refine the selection.

We provide locks and mouseover selection on polyline segments to convey interesting valid combinations of component alternatives. The atomic unit of an interaction is a system alternative $(\mathbf{v}_i, \mathbf{v}_j) \in V_{C_i} \times V_{C_j}$, i.e., a one-to-one combination of component alternatives. Any interaction taking place on one part of a system alternative also applies to the rest of the system alternative. Due to the combinatorial nature, multiple system alternatives might pass through the same polyline segment, in particular where shared axes are adjacent. An interaction with a polyline segment can thus lead to more than one system alternative being hit.

To specify the second selector, a filtered polyline segment can be clicked to lock the associated set of combinations, updating the selection to the respective subset of the filtered alternatives (Figure 7.8c). Only one segment can be locked at a time. A lock is active until it is unlocked (by clicking again) or moved to another polyline segment (by clicking the respective segment, see Figure 7.8e). Unlocking a lock makes the selection fall back to the superset of filtered alternatives.

To specify the third selector, the selection resulting from the lock can be refined via mouseover (Figure 7.8d). If there is no active lock, the mouseover operates on the set of filtered alternatives (Figure 7.8b). The mouseover interaction is temporary: when the cursor leaves the hovered polyline segment, the selection falls back to the set of locked alternatives or to the set of filtered alternatives if no lock is active.

Anything that remains in the selection after applying the current cascade of filters, lock, and mouseover is highlighted. At the end of an analysis, this is usually a unique combination of component alternatives, i.e., the final system choice (Figure 7.8f).

7.4 VALIDATION

Composite Parallel Coordinates provide a novel approach to a simultaneous exploration and analysis of multiple interacting data sets. To validate its domain usefulness in terms of problem-solving characteristics, we report on two usage scenarios and one case study [113] from distinct application domains. In the case study, we particularly reflect on observing an analysis conversation between two engineering experts. The results suggest that our technique supports the identified analysis tasks for making co-dependent multi-attribute choices.

These three real-world scenarios showcase how COMPO*SED helps users simultaneously explore linked component ensembles for the analysis of complex systems. The data sets exhibit different properties regarding unique, shared, and related design parameters and criteria. In all three scenarios, the visualization enabled users to apply constraints and observe their combined effects on both the component and the system level. In particular, it supported decision-makers in investigating how a component-specific choice affects the system performance and the availability of interoperable component alternatives.

7.4.1 *Usage Scenario I: Power Plant Operation*

District heating describes a method for delivering space heating or hot water to buildings via insulated pipe systems. Power plants that serve as heat suppliers mostly fire fossil fuel, biomass, or waste. This usage scenario stems from a long-term collaboration with engineers who are responsible for the operation of a district heating power plant. The power plant is located in Austria, and the data has been anonymized to not reveal ownership. The power plant consists of two neighboring blocks to burn different types of combustion material. One block uses water, while the other block uses air as a cooling material.

Rather than interoperability, it is the joint production of heat that prohibits an optimization of the blocks in an isolated manner. Operators of the power plant can use Composite Parallel Coordinates to decide which block to use to which extent in order to jointly produce a certain amount of heat requested by the operation plan. If the outcome of one block decreases, the outcome of the other block has to be increased. This is typically based on the domain knowledge of operation engineers. Using data analysis for power plant operation is still a novel approach. The engineers only have basic experience in data analysis and data visualization. With the Composite Parallel Coordinates, they can, for the first time, study 1) how environmental conditions influence the possible operation modes of the power plant and 2) how the parameters of one block influence the operation of the second block. The data was generated using a simulation model.

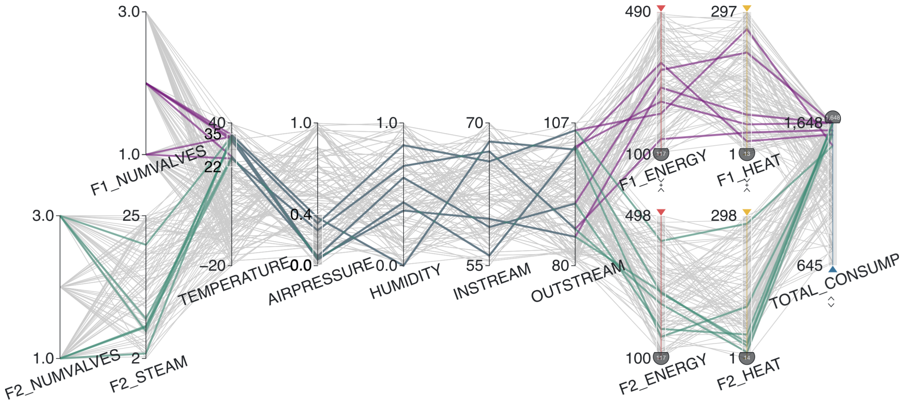

Figure 7.9: Power plant operation: a certain amount of heat is to be produced by a combination of blocks (top and bottom) that share the same environment (center). High temperature and low air pressure lead to a high overall consumption of combustion material (right).

The heat produced by each block depends on various factors. Block-specific design parameters include combustion material, cooling parameters, and the number of active valves. Environmental conditions like temperature, humidity, and air pressure are shared by both blocks because the blocks are equally affected by their changes. The different consumptions and efficiencies associated with the produced heat in each block are related and can be composed into system criteria.

Engineers need to constantly regulate the operation of the blocks during the day. The main trade-off lies in producing a maximum outcome while consuming the least possible amount of combustion material. The goal is to distribute the production of the requested heat to both blocks such that the yield, i.e., the difference between the price for heat on the market and the operating costs, is maximized.

Every polyline in the plot corresponds to one possible operation mode of the power plant (Figure 7.9). The two blocks are shown as separate pathways. First, engineers can study the influence of the shared environmental parameters on the operation of the power plant. Merging the blocks' individual consumptions via addition, they can observe that high outside temperatures and low air pressure both lead to higher overall consumption of combustion material and therefore high costs. At the component level, only a low number of valves is needed for the air-cooled block to reach a high outcome when outside temperatures are high. An analysis of the separated variable spaces of both blocks shows how the number of valves of one block influences the operation of the other block. With two active valves for the water-cooled block (since this ensures a low consumption of combustion material), a similar combustion-saving setting for the air-cooled block relies on high air pressure and high air humidity – thus, it highly depends on non-controllable environmental conditions.

The Composite Parallel Coordinates enable operators to see all involved parameters at a glance (**T_1** Overview) and to understand the dependencies between parameter settings of different blocks (**T_6** Navigation). With the novel representation, engineers are able to study the effect of outside temperatures on the needs of combustion material simultaneously for different blocks (**T_5** Connectivity). They are also able to see how a block-specific decision for a number of valves affects the energy production in the other block (**T_8** Partial Choice). In this sense, Composite Parallel Coordinates open up new possibilities for investigating multi-block power plants.

7.4.2 *Case Study: Magnetic-Geared Motor Design*

Magnetic-geared motors (MGM) are suited for industrial applications that require high power densities, e.g., wind energy or ship propulsion. To achieve the desired outcome, the driving motor and gear need to interact effectively. This case study was informed by an observational study, where two engineers collaboratively worked with the tool. We report on their qualitative feedback at the end of this section. The one-and-a-half-hour remote study was recorded. Both engineers have a background in mechatronics and years of experience in electric drive design. Their daily work involves complex simulations and optimizations of geometries, magnetics, thermal conditions, and their interplay. They are familiar with basic visualizations and brushing techniques, including standard parallel coordinates. We primarily wanted to identify aspects of our technique that are particularly relevant to the engineers' decision-making. Thus, we did not impose a pre-defined setting but emphasized free discussions on one of their real-world design optimization use cases. We only prescribed the high-level task to analyze the data and choose the most preferred motor-gear combination. After a brief introduction to the functionality of the tool, the engineers explored the Composite Parallel Coordinates on their own.

Data Analysis

Our domain experts used COMPO*SED to analyze 392 combinations of motor and gear alternatives. The data result from optimizations they conducted to investigate a side-by-side arrangement of motor and gear [161]. The design parameters in each data set represent geometric properties and operating conditions. The gear ratio (*G12*) and output specifications (*N2*, *P2*, *T2*) are common to both data sets (Figure 7.1a). Related criteria that might be assembled into system criteria are component lengths (*LFE*), power losses (*PV*), and efficiencies (*ETA*).

The experts aim at parameter settings of motor and gear that lead to a high overall efficiency with a low construction volume and torque ripple. First, the experts observe how the motor and gear alternatives distribute along the unique and connected parts of the system (**T_1**

Overview). Viewing the separate pathways, they notice a motor outlier with drastically lower efficiency (*ETA_MOT*) than all other motor alternatives. They also recognize two clusters of gear alternatives that significantly differ with respect to efficiency (*ETA_MG*) and torque ripple (*T1_RIPPLE*).

The primary objective is the efficiency of the entire system. At the motor level, they filter out the outlier with low efficiency (**T3** Filter). Next, they merge the efficiencies of motor and gear using multiplication (Figure 7.1b, *ETA_MG * ETA_MOT*). On the collapsed axis, they restrict the overall efficiency to high values (**T3** Filter, Figure 7.1d). From the separate pathways, the experts notice that, to their favor, motors with high current density (*JS*) and high copper losses (*PCU*) are excluded.

The secondary objective is the length as an approximation of the construction volume. It needs to be filtered at the component level. Otherwise, the engineers could not recognize undesired combinations where the total length is acceptable, but motor and gear lengths differ significantly. Restricting the motor and gear alternatives to small lengths each (**T3** Filter, Figure 7.1c, *LFE_MOT* and *LFE_MG*), the experts notice that gear alternatives with preferred low torque ripple are not in the selection anymore. Undoing the previous filter actions one by one reveals that the previous system efficiency maximization excluded them. This correlation was not known before. The engineers expect it to originate from problem-specific boundary conditions of the optimization.

The experts now face a conflict between a system-level criterion and a gear criterion (**T2** Competition). They decide to not sacrifice the gear criterion too early and rather investigate the trade-off from the reverse perspective. Brushing the cluster of gear alternatives with low values for their unique parameter torque ripple (**T3** Filter, Figure 7.1c, *T1_RIPPLE*) leaves the engineers with about 50 MGMs still offering acceptable system efficiencies (**T6** Navigation, **T8** Partial Choice).

The current selection is associated with short gears. From their experience, the engineers anticipate that this could induce less output power (*P2*) of the system (**T2** Competition). However, the output power should not be too low. Brushing the upper half of the respective shared axis results in a dozen motor-gear combinations (**T3** Filter). Merging the components' length axes via addition (*LFE_MG + LFE_MOT*) reveals that the selection still contains short MGMs (**T6** Navigation). Other properties offer potential for further drill-down.

Two clusters of gear alternatives can be observed for the unique parameter flux density (*OBS_AIR*): one with higher flux density and one with low flux density. Brushing the latter results in six selected MGMs, which still cover a wide range of system efficiencies (**T8** Partial Choice). One outlier with significantly higher total length is excluded.

The remaining five magnetic-geared motors are on par with respect to their performance. Manufacturing benefits are thus pivotal. If gear magnets are nearly squared, their mounting direction might get mixed up, leading to wrong magnetization. If their distance is too low, they are difficult to mount. After a detailed comparison (**T4** Subjective Evaluation), the MGM design with the largest gear magnet rectangularity (*D_BM_HM*) and distance (*D_PM3*) is chosen (Figure 7.1e).

COMPO*SED allowed for constant switching between the component level and system level and between overview and detail. Unlike before, the experts did not have to go back and forth between individual component optimizations. Instead, the component alternatives and their dependencies could be explored simultaneously. At both levels, design parameters and criteria could be equivalently used for real-time filtering. This helped the engineers directly take into account the effects of each component decision on the system operation. Rather than choosing the first working solution, the engineers could learn which combinations and what level of performance were achievable under which conditions.

Expert Conversation

Although field observations and think-aloud walkthroughs are common evaluation methods, performing them with pairs of domain experts (E1 and E2) is rather rare [138]. Our motives were slightly different from those of studies in computer-supported cooperative work [188] and collaborative visualization [111], which primarily aim to assess teamwork. First, the conversation resembled the day-to-day practice of our experts, who analyze and discuss complex optimization problems collaboratively. Second, we hypothesized that a natural conversation between like-minded colleagues yields more valuable insights than an artificial monologue of a single expert.

We found that the overview of all involved design parameters and criteria – in particular their different roles – is the primary advantage of Composite Parallel Coordinates: *"A lot of information is presented in a clear and compact way"* (E1). They also adequately support filtering both at the component level and the system level: *"If you drag the slider slowly, you can easily trace which alternatives drop out and at which point they join back in"* (E2). In fact, the ability to view and constrain individual components while also observing system-level properties was perceived as a significant advantage: *"If you would restrict the system length, e.g., to 70 mm, you might end up with a 60 mm motor and a 10 mm gear, which would simply be useless"* (E1). Although their routine involves making complex decisions collaboratively, multiple users interacting in realtime with the same visualization is not desired, as they can no longer understand what has led to the final outcome: *"The collaborative decision-making is about considering the next steps together, not about speeding up the interaction"* (E1). COMPO*SED did not directly reduce the time required for a choice, but the experts reported that it helps

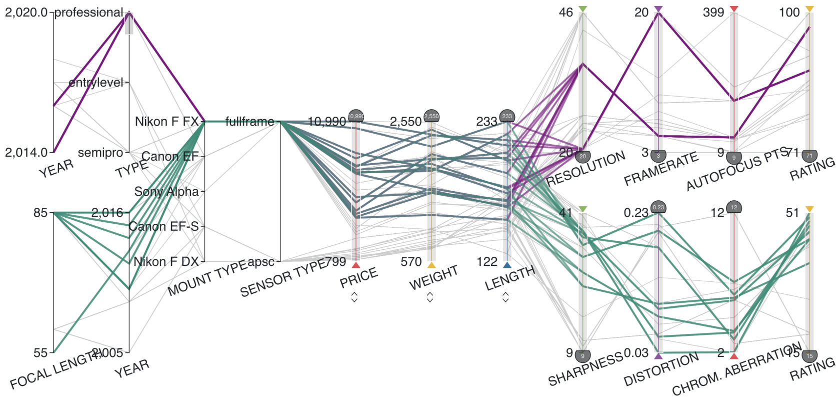

Figure 7.10: A photographer can choose from combinations of three cameras (purple) and 13 lenses (green). Only components with the same mount type are compatible (third from left). Component prices, weights, and lengths are added up to system criteria (center).

avoid optimization iterations. With existing tools, their choice is based on a subset of the most important parameters. If inconsistencies arise during validation, they enter an additional iteration. Such iterations and the additional time are avoided by the more comprehensive picture our technique offers: *"With COMPO*SED, we can keep an eye on all parameters right from the beginning"* (E1).

The collaborative analysis session was highly similar to pair-programming: one expert, the driver, interacted with the visualization, while the other, the navigator, kept an eye on effects and hinted at aspects to address. The conversation was free-flowing, interrupted only by considerations of what to look at next. The engineers communicated by agreeing upon next steps, refining each other's explanations, and at times even correcting each other. They also drew the other's attention toward interesting regions in the visualization. Watching the engineers learn not only from the data but also from each other provided us with insights that we might not have gained otherwise.

7.4.3 *Usage Scenario II: Camera-Lens Purchase*

The previous scenarios dealt with systems where one component alternative was exclusively compatible with exactly one alternative of another component. However, some system design problems might require the component alternatives to be combined more freely. An example are cameras and lenses, where one camera body can be equipped with different lenses and one lens can be mounted on multiple camera bodies. Aiming at optimized equipment, a photographer might use COMPO*SED to decide whether to buy a lens to mount on her semi-professional camera body or upgrade to a professional camera body requiring new lenses due to a different mount type.

The former case requires the analysis of a one-to-many relationship. Brushing her existing camera body, the photographer compares the five compatible lenses (**T5** Connectivity). They cover a wide range of prices, weights, and lengths (**T8** Partial Choice) while exhibiting similar ratings. The photographer excludes two lenses that are located towards the upper ends of the price, weight, and length ranges while not performing exclusively better in the remaining criteria. The remaining three lenses have similar prices. The final choice for one of them cannot be made at the component level. Instead, the photographer needs to consider how their characteristics like sharpness, distortion, etc. work together with the existing camera's resolution, framerate, and so on.

To further improve the performance, it might be beneficial to replace the camera body with a professional one (**T9** Alternative Replacement). This requires the analysis of a many-to-many relationship. Brushing the professional camera bodies, the photographer is left with combinations of three camera bodies and 13 compatible lenses (Figure 7.10). To not miss a preferable combination, the photographer first looks for alternatives that yield a good performance while being compatible with many other alternatives (**T7** Key Component Alternatives). Applying her total budget as a constraint, the camera body with the highest rating and a high resolution remains. It is still compatible with six of the 14 lenses, leaving enough room to further exploit the optimization potential at the lens level.

7.5 DISCUSSION AND LIMITATIONS

Parallel coordinates are certainly one of the more complex visualization techniques. In the previous design study (Chapter 6), we learned that conventional parallel coordinates are accessible for analysts performing single multi-attribute choices. Nevertheless, the question remains whether the added complexity of the novel parallel coordinates variant matches its increased usage value.

Composite Parallel Coordinates are not merely two linked visualizations. A distinctive contribution of our technique is the possibility to jointly explore the alternatives of multiple components, their combinations, and constraints. The relations between system and component properties can be perceived from the side-by-side layout and the interaction with composite axes. In particular, our approach shows one-to-many and many-to-many combinations of matching alternatives explicitly by extending the idea of linked axes across components. Our strategy follows the recommendation to integrate views with an explicit linking when relations between items of different data sets are of particular importance [115]. The cost of added visual complexity is mitigated by filters, which our experts perceived as a powerful tool. In contrast, the scenarios did not require the parallelogram strategy to avoid line crossings. The numbers of polylines seem to have

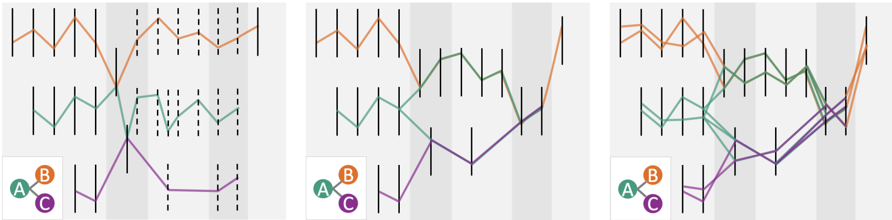

Figure 7.11: A sample case with three components where A interacts with B and C. On the left, the collapsable axes are dotted. In the center, all composite axes are collapsed. On the right, three complete system alternatives shown.

been manageable, which leaves us with the open question at which point the strategy develops its full potential. A number of techniques have been proposed to address different issues associated with dense line charts, including density estimation [95], edge bundling [310], and importance-driven blending order [273]. These techniques can be transferred to Composite Parallel Coordinates, potentially involving a particular treatment of component connections. To what extent the explicit but visually complex depiction of many-to-many relations simplifies the analysis is yet to be examined.

We have demonstrated the working principle for two components. A major limitation of the composite layout is that it does not effortlessly scale to an arbitrary number of data sets. While the layout is generally open to stacking multiple parallel coordinates plots, any plot can only be directly connected to two neighbors above and below (see Figure 7.11). The vertical positioning according to component links is remarkably similar to the axis-ordering problem of conventional parallel coordinates. We hypothesize that existing solutions (e.g., linearization of node-link representations, interactive reordering, or aggregation) can be adapted to overcome this limitation on the vertical axis. Stacking more than two components leads to more complex branching of the polylines, requiring dedicated rendering and interaction techniques to trace individual polylines. To what extent high polyline density can be mitigated by filtering is yet to be examined. Showing relations between non-adjacent components will likely introduce additional visual clutter. Additional ranges on the horizontal axes are needed to depict related attributes of non-adjacent components (see the second gray area in Figure 7.11, where parameters of all three components are totalized). With a certain number of components making up the system structure, a dedicated navigation strategy (e.g., using a minimap) might generally be required.

Another limitation is that composite axes represent only one-to-one mappings of related properties from different components, with rather simple predicates and aggregations. Where other constellations are required, e.g., one-to-many mappings of related properties, a more flexible representation of composite axes is needed.

The validation scenarios showcase different relations between components that can be explored using COMPO*SED. Two of three applications have only been illustrated through a usage scenario rather than being verified with a case study. Still, all three validation scenarios indicate that our technique adapts to different types of data and tasks, providing an effective means for co-dependent multi-attribute choices.

7.6 CONCLUSION

This chapter investigated the demands that are posed on a visualization technique to effectively move from supporting single multi-attribute choices to supporting two or more choices that are co-dependent. Co-dependent choices introduce unprecedented challenges. The essential challenges are 1) the huge combinatorial solution space, 2) side-effects that prevent decisions from going well together, and 3) emergent properties that the individual choices do not possess when made separately. Visualizations need to provide an exploration of options associated with the same choice but also reveal the performance of option combinations across choices.

This chapter presented COMPO*SED, a novel parallel coordinates variant for the visualization of distinct but related data sets to help humans make decisions where choices affect each other. Each data set contains the options belonging to the same choice. Side-effects in choosing the most preferred combination of options require trade-offs beyond one single choice. Parallel coordinates plots depicting each data set are juxtaposed. A vertical stacking of the plots has turned out to be the most suitable layout for communicating dependencies between the choices. Where attributes of different choices share a similar semantic meaning, the associated axes can be merged. This effectively visualizes the side-effects between intermediate decisions as the bottleneck of co-dependent choices. Three cascading interaction patterns enable analysts to explore option combinations under consideration of individual option properties and combined system properties. In particular, temporarily locking a subset of options is needed to interactively refine the selection in a juxtaposed plot. Qualitative feedback from three real-world scenarios and an observed expert conversation underlines the need for co-dependent choice support and suggests the effectiveness of our visualization technique for this purpose.

An obvious limitation of the composite layout is that it does not effortlessly scale to an arbitrary number of choices involved. While the vertical layout is open to stacking more than two parallel coordinates plots, future layout and interaction considerations primarily need to address the relations between non-adjacent plots. While the qualitative feedback was positive, the trade-off between visual complexity and usefulness should be investigated in more detail, in particular when targeting more than two co-dependent choices.

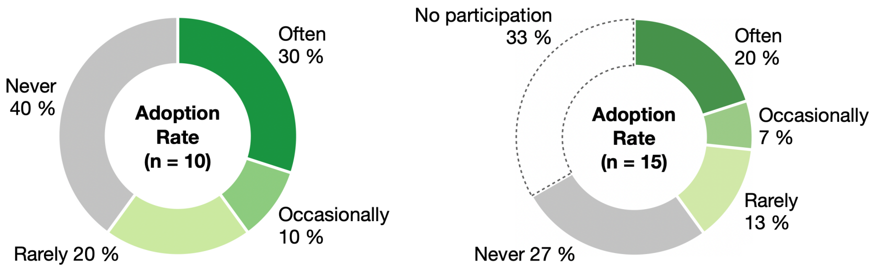

(a) Study participants (b) Entire target group

Figure 8.1: **RQ$_1$** Adoption Rate: (a) 60% of the study participants tried out PAVED on their own initiative. (b) If all potential target users are taken as a reference, the rate of known adopters is 40%.

8 LONG-TERM EVALUATION OF DECISION SUPPORT

In Chapter 6, we described the development of PAVED, an interactive parallel coordinates visualization to support choice-making in motor design. At the end of our design study [54], PAVED's short-term usefulness was positively evaluated in a field study. The domain experts then decided to integrate PAVED into their daily motor design workflows. More precisely, PAVED was made accessible from the interface of their in-house tool to explore the results of an optimization project. This can be considered a first success.

However, the decision was made during a period of intense collaboration and Brehmer et al. observed that tool use typically decreases after such a period [37]. Consequently, our research question was: after four years, is PAVED still being used in the experts' daily work? A permanent adoption in the sense of repeated, self-initiated use for the engineers' daily activities would provide evidence that PAVED in fact addresses the true needs of the target users. In the domain of motor design, studying adoption might be particularly meaningful, because motor designers are free to choose any tool they consider useful in a project. Despite the importance and meaningfulness of long-term adoption, Section 3.1.3 showed that few visualization (design study) papers follow up and report on the long-term usage of their proposed tools. In fact, no publication reported to have repeated the same questionnaire after a certain period of time.

For our design study, we close this gap by revisiting the usefulness of our tool in the targeted application domain four years after the initial summative evaluation. The results provide a comprehensive picture of our tool's adoption readiness level. They revealed a small number of satisfied power-users, who regularly benefit from PAVED in their daily work and *"would not want to miss it"* (A2). These users also reconfirmed the high usability of PAVED. Only one target user purposefully abandoned PAVED due to an integration issue. The responses of 10 participants in total suggest a general adoption rate of 60% in the four years, where domain experts independently tried out PAVED for their day-to-day tasks in motor design (Figure 8.1).

What the reader can expect from this chapter:

- The first long-term study of decision support that investigates the self-initiated use of PAVED after four years (Sec. 8.1 and 8.2).
- An analysis and discussion of the data showing that the experts still routinely use the tool for its lossless overview of all options and reduced interaction mechanisms (Sec. 8.3 and 8.4).

8.1 AIM OF THE STUDY

Our objective is to evaluate the day-to-day usage of PAVED and derive insights about its long-term benefits for the engineers' tasks. For this purpose, we set up a post-deployment study using an online questionnaire. Again, we collected data for a qualitative analysis, which is augmented with a quantitative usability scale where applicable. Contrary to the previous short-term evaluation, however, the researcher did not engage in an observation or interview.

The goal is to validate the domain usefulness of PAVED with a particular focus on adoption and routine use. Of the adopters, we wanted to know for which tasks they used PAVED, what functionality was (not) useful, what challenges they faced, and how well the tool is integrated with their existing workflows. We also aimed at a quantitative assessment of the long-term usability. Of those who had not adopted PAVED, we wanted to learn about the reasons. If their tasks did not require to use PAVED, the failed adoption was not about the tool itself. But if they actively decided against its use, we wanted to know why.

Our research questions can be summarized as follows:

RQ1 **Adoption Rate** – What portion of target users did adopt PAVED for regular use as part of their daily work? Who are they?

RQ2 **Usage Context** – What are the circumstances, under which PAVED is still used in the application domain? Did they change?

RQ3 **Usability** – How does the perceived usability compare to that four years ago?

RQ4 **Reasons for Refusal** – Why did some target users not adopt PAVED?

8.2 METHODOLOGY

In the context of our collaboration, the target users of PAVED are mechatronic engineers who are involved in the design of electric motors using their in-house optimization tool *SyMSpace* [248]. From our primary domain expert, we received a list of 15 engineers to whom this applies. Some of them have already participated in the evaluation four years ago (Section 6.4). Among this group, we are interested in the adoption rate following our design study and in the circumstances, under which PAVED is (not) used.

For this purpose, we collected the engineers' experiences with PAVED based on their voluntary, self-initiated use since the deployment four years ago. With the approval of our primary domain expert but without any further announcement, we invited the target users via e-mail. The invitation briefly described the research objective and provided a non-personalized link to the online questionnaire, which we hosted on a local installation of SoSci Survey at TU Darmstadt. The questionnaire was open over a time period of two weeks. Sending out a friendly reminder after one week and again on the last day of the survey period was quite effective. By inviting all potential target users regardless of their actual usage behavior, we avoided survivorship bias and retained the opportunity to derive insights also from supposedly failed adoption cases.

The aim of our study was purely academic and did not involve any expected consequences for the target users that could be framed as an incentive to participate. To minimize the refusal rate, it was thus important to develop a practical and time-effective yet meaningful questionnaire. We aimed at a time commitment of 10 to 15 minutes and closed-ended questions wherever appropriate. We compiled the questionnaire by drawing inspiration from sample questions provided by other visualization researchers to assess the adoption of a visualization. Lam et al. propose sample questions addressing two different perspectives: 1) focusing on the intended work environment by studying the workflows and practices with a visualization and 2) focusing on the specific visualization by studying users' personal (long-term) experience with it [138]. In assessing the adoption of their tool *Jigsaw*, Kang and Stasko focus on usage purpose, comparison to traditional methods, (not) useful and missing features, and barriers [121]. Their sample questions largely overlap with those provided by Lam and colleagues. Brehmer et al. conclude their design study of *Overview* with an assessment of who adopted their tool, how it had been used, whether it is still in use, and what problems users reported [37].

On this collection of sample questions, we performed a thematic analysis to compile a set of important aspects to cover in our questionnaire. In an inductive coding manner, we labeled the topics in the sample questions and grouped recurring codes into five themes that relate to our research questions:

- **Demography** – What responsibilities (non-)adopters have and for how long they have been employed in their job role (**RQ1**)
- **Usage Behavior** – For what daily activities and data PAVED has been used, whether it is still in use, and how important and recommendable it is for the daily work (**RQ2**)
- **Features** – What functionalities did (not) work well or were missing, what challenges were encountered, and how PAVED compares to existing methods (**RQ2**)
- **Integration** – How well PAVED blended into the domain experts' existing workflows in terms of access, data handover, and functionality coordination (**RQ2**)
- **Hindering Barriers** – What limitations hindered the adoption of PAVED (**RQ4**)

To assess how the perceived usability has evolved in the past four years (**RQ3**), we additionally included **System Usability** as a theme. We replicated the use of the System Usability Scale (SUS) from our design study, where ten statements are rated on a 5-point Likert scale [232]. For the five remaining themes, we selected a handful of representative questions each.

We started the questionnaire by asking the participants whether they were aware of the possibility to use PAVED before taking part in the questionnaire and, if yes, how regularly they used PAVED. Only those who reported to have used PAVED at least rarely were asked in detail about their experience by working through the above themes. Figure 8.2 depicts how participants are directed along different paths in the questionnaire depending on their answers. We closed the questionnaire by asking the participants whether there was anything else they would like to tell us about the role of PAVED for their daily work. In total, we ended up with 31 questions, of which 21 were closed-ended (yes/no and rating scales). The ten open-ended questions were text inputs, either as multi-line input fields or as free mentions where each answer is stored individually. Four open-ended questions were optional. Assuming an average of 25 seconds completion time per question [143], this matches our initial ambition to require a maximum of 10 to 15 minutes time commitment.

The wording and obligation of questions as well as the questionnaire navigation have been validated by a small-scale preliminary study, for which five research fellows were recruited. The collected comments addressed the clarification of intentions and the scope of questions as well as layout issues like the placement of consent and contact details or the labeling of rating scales. Subject to a detailed discussion have been the questions whether participants should be informed that a particular answer will finish the questionnaire and whether questions should include examples of what kind of answers is expected.

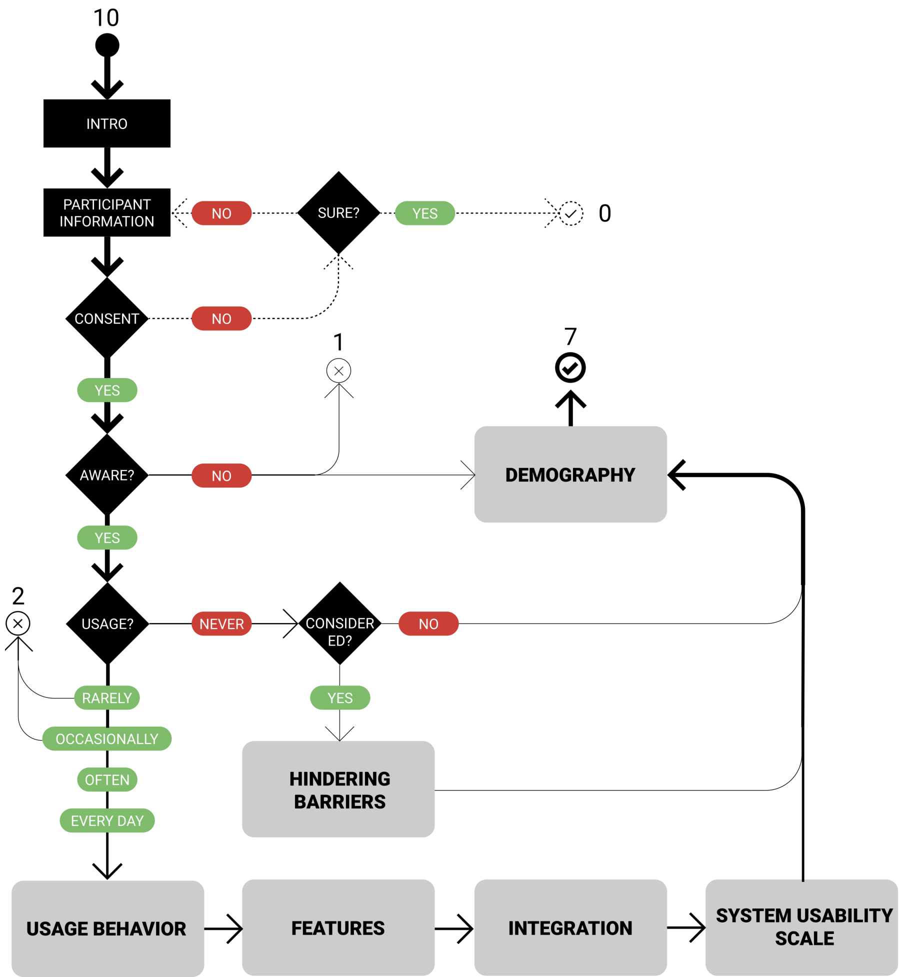

Figure 8.2: Flow chart of the logic that directs participants along different routes within the questionnaire. Four adopters traversed the outermost loop covering details of the usage behavior as well as perception of features, integration aspects, and system usability. Three participants were guided via a shorter route because they reported to have not used PAVED. Another three participants dropped out in the middle of the questionnaire ⊗.

8.3 RESULTS

Out of the 15 engineers in our target group, 10 responded to our invitation to participate in the study. This is above the average online survey response rate of 44.1% [299]. Figure 8.2 depicts the flow of respondents through the different branches of the questionnaire depending on their answers. Their completion times varied accordingly. They ranged from 30 seconds to five minutes for participants who were guided via the short route, while participants who answered the detailed theme blocks spent 10:30 minutes over 10:50 and 22:30 to 51:30 minutes. In the following, we analyze the collected data with respect to our four research questions.

RQ1 Adoption Rate

We define adoption as *"Did a [user] freely choose the tool for their own investigation, rather than trying out the tool in response to direct solicitation by the researchers?"* [37]. According to this definition, adoption occurred in six of the ten response cases that we report here (Figure 8.1a). These six participants reported to have used PAVED rarely to often but not daily in the last four years. The remaining four participants reported to have never used PAVED. The reasons for this will be presented in the respective paragraph.

Of the six adopters, three are still using PAVED for their work today. Two of them left the questionnaire immediately after they stated how often they used PAVED such that we do not know about their on-going use. Finally, one of them reported to not currently use the tool as part of the day-to-day work. The reason for this might be found in the participant's job responsibilities as a manager. In this position, the role of PAVED is to demonstrate the high-level potential of optimization software in technical customer discussions rather than being used for the actual motor design.

Five, or one third, of the invited target users did not participate in the study, such that we cannot make any reliable statements about their adoption of PAVED (Figure 8.1b). From their auto-replies, we learned that two of these five invitees were on parental leave at the time of the study. We know that one of them has in fact used PAVED one and a half years after the deployment based on questions and feedback we received via e-mail.

RQ2 Usage Context

Those four adopters of PAVED who did not drop out of the questionnaire answered the questions that detail the circumstances, under which PAVED is still used in the application domain: usage behavior, features, and integration. Of the four participants, three reported to have used PAVED often in their daily work, while the remaining participant used it rarely. We characterize the current usage context of PAVED based on the participants' ratings of the closed-ended questions (Table 8.1) as well as a coding and sorting of their responses to the open-ended questions. In the following, we will describe our insights with respect to each of the themes.

In line with the task it was designed for (Section 6.1.3), PAVED was mainly used for exploring motor designs and choosing an optimal design (A1, A2, A3). An integral part of this was to analyze the relations between the different properties (A2) and to communicate what performance is achievable under which conditions to the customer (A3). Revisiting decisions with customers was already identified as a hot topic in the first evaluation (Section 6.4.2). Rather than using it for choice-making in an actual motor design project, one participant used PAVED with a demo data set in customer meetings to showcase the

	Usage Behavior			Features	Integration			
	Importance for daily work	Still used today	Recommend	Comparison to traditional tools	Access from within SyMSpace	Handover of data and filters	Coordination of functionalities	Blend into workflow
A1	●●●●○	●	●	●●●●○	●●●●○	●●●●○	●●●●○	●●○○○
A2	●●●●●	●	●	●●●●●	●●●○○	●●●●○	●●●●○	●●●●●
A3	●●●●●	●	●	●●●○○	●●●●●	●●●●○	●●●●●	●●●●●
A4	●●●○○	○	●	●●●●○	●●●●○	●●●●●	●●●●●	●●●●○

Table 8.1: **RQ2** Usage Context: ratings of the closed-ended questions targeting the usage context grouped by themes. Together with the qualitative feedback, the responses indicate that PAVED addresses the true needs of the target users.

general potential of visualization and optimization for the design of electric motors (A4). The other three participants, who still use PAVED today, expressed that the tool is *"very important"* for their daily work (Table 8.1, Usage Behavior, average: 4.25/5, range: 3-5). Unfortunately, the participants did not understand what we meant when asking for the characteristics and semantics of the explored data sets, such that we cannot draw conclusions about the diversity regarding the number of attributes and alternatives or the subject to be decided upon. All four participants, however, would recommend the tool to a new colleague in a similar job role (Table 8.1, Usage Behavior). This aligns with the high usability we observed, where tools with a score above 80.3 are more likely to be recommended [232].

PAVED's primary view, the parallel coordinates visualization, was mentioned as most useful (A4), especially for exploring the dependencies and potential conflicts between attributes (A2). One participant also appreciated the list of options with their raw attribute values shown in the table view (A2) and thereby, once again, confirmed the relevance of tabular visualizations for decision-making (Section 6.5.2). The participants proposed to extend the parallel coordinates visualization to view additional result sets (A3), for example to compare multiple Pareto fronts (A1). Being able to show and hide axes flexibly was perceived as particularly helpful (A2). This feedback was supplemented with the suggestion to introduce a semantic grouping of available attributes to facilitate multi-selection (A2). Also, the tool was experienced to run *"quite slow"* when many attributes are shown (A1). In line with our view that interaction is essential for an effective use of parallel coordinates (Section 6.3.3), the participants also highlighted the dynamic filtering as particularly useful (A2, A4). Contrary to the first evaluation (Section 6.4.2), our proposed preference brushes for optimization criteria do not seem to have provoked comprehensibility issues this time. Instead, one participant found the restriction of design parameters *"not very intuitive"* (A1). As this slightly conflicts with our assumption (and learning from the first evaluation) that the target users were familiar with standard range brushes, this issue would be interesting to further investigate. Overall, the participants rated PAVED as *"somewhat more"* helpful for their work than traditional approaches (Table 8.1, Features, average: 4/5, range: 3-5).

Given the participants' ratings (Table 8.1, Integration), integration issues between PAVED and the engineers' in-house tool SyMSpace are not likely a potential reason for failed adoption. The responses to the four integration aspects indicate that PAVED is used in a well-functioning symbiosis with SyMSpace. The participants found the PAVED interface *"easy"* to open from within a SyMSpace workflow (average: 4/5, range: 3-5) and were *"very satisfied"* with the handover of data and selections between both tools (average: 4.25/5, range: 4-5). They also agreed that PAVED and SyMSpace are *"complementing"* each other in terms of functionalities (average: 4.5/5, range: 4-5). The strongest disagreement among participants occurred on how well PAVED blends into the SyMSpace analysis workflows (average: 4/5, range: 2-5), where one participant said *"poorly"* (A1) and two said *"very well"* (A2, A3).

RQ3 Usability

The same four adopters filled out the quantitative System Usability Scale, which we used to compare the perceived usability of PAVED to the previous assessment from our design study. While we hoped to confirm the previously high usability, we hypothesized that the self-initiated usage might reveal usability issues that did not occur in the observed walk-through sessions four years ago. Although the scores of individual SUS dimensions are not meaningful on their own [39], we were also interested whether the responses to the individual statements differ from those we received four years ago.

The adopters' ratings suggest that PAVED still provides an *excellent* usability, according to the adjective equivalent of the achieved SUS score [18]. Based on four participants, it achieved a score of 86.9 out of 100, which is highly above the average score of 68 [232]. We present the individual scores broken down by question in Table 8.2. We noticed that PAVED scored highest on consistency (Q6), where all participants agreed on the strongest possible approval. The dimension ease of use presents a particularly interesting case. While it also received the highest possible score for the negatively formulated statement (Q8, i.e., *"I found PAVED very cumbersome to use."*), its positively formulated complement (Q3, i.e., *"I thought PAVED was easy to use."*) scored lower. Four years ago, the average ratings of both statements differed only marginally (9.7 versus 10 out of 10). In the present case, the average

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Total
A1	10	7.5	7.5	7.5	7.5	10	7.5	10	7.5	10	85
A2	7.5	7.5	7.5	7.5	7.5	10	7.5	10	7.5	7.5	80
A3	10	10	10	10	10	10	10	10	10	10	100
A4	7.5	10	5	7.5	7.5	10	10	10	7.5	7.5	82.5
Avg 2023	8.8	8.8	7.5	8.1	8.1	10	8.8	10	8.1	8.8	**86.9**
Avg 2019	7.8	9.4	9.7	8.8	8.1	8.1	9.4	10	8.4	9.7	**89.4**

Table 8.2: **RQ3** Usability: results of the long-term System Usability Scale [232]. The tool achieved a total score of 86.9 out of 100 (n = 4) compared to 89.4 (n = 8) four years ago (Section 6.4.2). Back then, interest in frequent use (Q1) received the lowest score, while ease of use (Q8) was rated particularly high. Four years later, consistency (Q6) scored highest, while the ratings for ease of use diverged (Q3, Q8).

values are strongly influenced by participant A4, who took a neutral position for the positive formulation but fully rejected the negative formulation. Another possible explanation for a general discrepancy in the ease of use dimension is that strongly disagreeing with a tool's suggested cumbersome nature might generally be easier than strongly agreeing to its suggested easiness to use.

With an overall score of 86.9 out of 100 (n = 4), the self-initiated and undirected use of PAVED during the adoption phase scored similarly to the requested and prescribed use directly after the tool's deployment four years ago, which received an overall score of 89.4 (n = 8). This indicates that the perceived usability of PAVED does not depend on the setting, in which it was obtained, i.e., a test-then-measure scenario with prior completion of tasks (after deployment) versus a retrospective evaluation without preceding tasks (adoption phase) [86]. Given the generally high level of scores received (i.e., within the range of 7.5 to 10 out of 10 for both assessments), the differences pointed out in the following comparison with the previous SUS assessment remain comparatively small. The interest in frequent use (Q1) initially received the lowest ratings with an agreement of 7.8 out of 10. This initial interest could be confirmed by a rate of 40% of the respondents who reported to have used PAVED sometimes or often. The interest in frequent use furthermore increased to a strong agreement score of 8.8 in the second evaluation. This indicates a generally positive trend among those who adopted PAVED. The dimensions that have received slightly decreased ratings compared to the first evaluation include tool complexity (Q2) and learnability, both for oneself (Q4 and Q10) and for others (Q7) [144]. In contrast, the consistency rating (Q6) raised from 8.1 in the first evaluation to the highest possible score of 10 in the present case. This might be attributed to our implementations of some suggestions that we received in the first evaluation.

RQ4 Reasons for Refusal

Out of the ten study participants, four had never used PAVED in the last four years. Two of them reported that, prior to taking part in the study, they had not known of the existence of PAVED to explore the optimization results of their in-house tool SyMSpace. One of them did not use the in-house tool for motor design where PAVED is integrated, which is an obvious reason to have never been in contact with PAVED. Technically, as a control design developer, this participant does not belong to our group of target users. The other unaware participant dropped out of the questionnaire immediately after saying so, such that we cannot know whether a similar reason applies here.

The remaining two participants reported that they were aware of PAVED but did not use it. One of them stated to not have considered its use at all. Only one out of four participants, a senior researcher with 18 years of experience, stated to have decided against using PAVED. The barrier in this case was an integration issue. The participant reported that it was unclear how to use PAVED within the user interface of the in-house optimization tool, from which PAVED could be accessed. To some extent, this contradicts our finding above, where integration issues were not existent among adopters.

8.4 SUMMARY AND DISCUSSION

In this section, we describe and discuss the results and logistics of conducting a post-deployment evaluation as a follow-up to a design study that was conducted four years ago.

We can report a final adoption rate of 60%, i.e., we know about six out of ten target users, who tried out PAVED on their own initiative. The reasons for an unsuccessful adoption of PAVED were three-fold. They ranged from not being aware of the possibility to use it over being aware but not having considered its use to actively deciding against its use for missing clarity regarding how to use it within the in-house software environment. Among those who regularly use PAVED for motor design, it is perceived as a *"very useful tool, especially for presentations"* (A3), which is well integrated with the surrounding domain software. In line with our primary design requirement (Section 6.5.1), the adopters agreed that the parallel coordinates provide a *"very fast overview of an optimization"* (A3) with dynamic filtering playing an important role. They applied the tool in the way we envisioned it to be used, suggesting that our domain characterization successfully informed the visual design. Further research could investigate to what extent our identified analysis tasks and decision process (Section 6.1.3) also hold true in other application scenarios. The high overall usability found in the initial assessment based on requested use (Section 6.4) could be reconfirmed based on self-initiated use four years later. The differences in the individual dimensions were not significant enough to suggest a severe need to take action in the short term, especially not in view of the few participants.

The low number of participants stems from the limited set of specific target users that we identified in advance, i.e., the motor designers at Linz Center of Mechatronics. While this helped us mitigate a selection bias that is inherent to many case studies [37], we could only analyze the responses of ten participants due to a refusal rate of 33% in an initially small group of 15 target users. Consequently, the absolute number of known PAVED adopters at this point remains fairly low.

A limitation of our study is the choice of an indirect online survey over direct observational methods for data collection. Contrary to previous works [37, 121], we sent out a questionnaire to collect the target users' experience with PAVED since the deployment. Our motivation was to reduce the time effort for both researchers and domain experts as we assumed the commitment to have decreased four years after collaboration. Given the moderate number of fifteen target users, observations and interviews would have been a feasible methodological choice and might have resulted in richer feedback. Another (additional) possibility for data collection would have been to include a usage tracking with the deployment. Independent of the actual research method, the comparability between the short-term and long-term evaluations is limited. Although both are conducted on choice tasks from the users' day-to-day work, the initial evaluation requested the tool use whereas the post-deployment evaluation involved voluntary, self-initiated use.

Although designed for motor design, PAVED is generalizable to choice tasks outside of the original target domain (Section 6.6). Further work is, however, needed to study the adoption of PAVED in a broader engineering context or even in other application domains where decision-making is essential. This requires promotion within the respective communities, such that potential users know of PAVED as a tool for making cost-benefit trade-offs.

The preliminary study raised the issue whether questions should include exemplary responses to ensure that participants understand what kind of answers the question is aiming for. On the other hand, participants might not describe their own (potentially different) experience but stick to the few given examples. For this reason, we moved from *"What data did you explore with PAVED (e.g., number of individuals, number of objectives, number of design parameters, ...)?"* to *"What are characteristics and semantics of the data sets you explored with PAVED?"* regarding the usage behavior. However, we received rather uniform and high-level answers in the direction of *"motor designs"*. One participant even apologized for not being more precise due to a lack of understanding where the question is heading. In retrospect, we should have decided to include exemplary answers or to more precisely outline expected details.

One potential reason for failed adoption is a mismatch between a participant's needs and those that the tool addresses. While this might be a result of misled task abstraction, it might also be a result of the participant not belonging to the targeted user group. The responses to our online survey were not rich enough to verify the correctness of our task abstractions, such that we need to rely on our previous confirmation (Section 6.4) in this respect. Instead, to obtain meaningful insights and a reliable adoption rate, it is important to determine whether a participant actually belongs to the identified group of target users. When designing the survey, researchers should carefully think about what exact information are needed to discriminate target users from other participants – and have the questions query all of these information. As we invited supposed target users (prefiltered by our primary domain expert), the questionnaire only discriminated participants based on their awareness and consideration of the possibility to use PAVED (see Figure 8.2). However, we have missed to query whether they were target users in the sense of SyMSpace users and motor designers at all. We noticed that this was not the case for all participants only from their self-reports. Because we could not ultimately determine the target users, we computed the adoption rate based on all valid responses, potentially resulting in an underestimation.

Similarly, statements about the regularity of tool usage need to be put in relation to the regularity with which participants face a task for which the tool was built. While rare tool usage might be a result of misled task abstraction, it might also be a result of the participant not facing a relevant task more often. We only verified this in the case of reported non-usage. For all other cases, we assumed the frame of reference to be daily confrontation with a relevant task. This might have resulted in an underestimation of the usage regularity.

8.5 CONCLUSION

Design studies that include an observation of their work's adoption and long-term usefulness in the target domain are rare. In this work, we add to the small number of longitudinal investigations by extending our design study (Chapter 6) on Pareto front visualization with a post-deployment study. Four years after the initial summative evaluation, we conducted an online survey with ten engineers to follow up on the role of PAVED for choice tasks in engineering design. We asked the participants about their usage behavior with PAVED, how they rate its usability, and for what reasons they possibly abandoned it. The aim of our study was to validate the long-term domain usefulness of PAVED against the background of its adoption and routine use.

The insights about how PAVED is used in practice are based on the valuable feedback of four participants who routinely used it to make or demonstrate choices about motor designs. Their qualitative responses show that, in their day-to-day practice, the participants appreciate PAVED for nothing less than it was designed for: a compact overview of all available options and attributes with reduced yet effective interaction mechanisms that help learn what performance is achievable under which conditions. Although our study focused on choices in engineering design, these findings can also inform the design of visualization tools targeting different decision tasks. For example, they add evidence to the successful application of parallel coordinates for the visualization of design options in various domains. They also support voices calling for thoroughly designed simple visual encodings rather than making design choices for novelty's sake.

Identifying factors that influence whether and how a tool is adopted in the intended work environment could help visualization researchers adjust their designs and collaborations accordingly. From our experience, we hypothesize that there is more to designing for permanent adoption than addressing the true needs of target users. Collaboration factors might include involving gatekeepers in the collaboration, discussing integration possibilities from early on, or raising awareness for the tool's existence through promotion activities. Understanding how domain experts earn praise for their work and aiming to support them in being successful might be another strategy towards adoption. In the case of our study, one aspect of this was that PAVED facilitated the communication with customers. Despite their benefits, interacting with visualizations requires time and cognitive effort. For a visualization to find its way into routine use, these costs should be minimized. For decision tasks, this can be achieved by considering proven design rationales like an efficient overview through lossless mapping or simple yet effective interaction mechanisms (Section 6.3.1).

A growing number of design studies propose visualization tools to support domain experts in their routine practice. Allowing some time for target users to adjust to a visualization support before assessing the usefulness and value of a tool in real-world working environments is advisable in this regard. It helps carve out the true needs of users, which might even require multiple iterations of design and deployment [37]. We thus argue that studying the self-initiated use of a proposed tool should be considered as a source of information in visualization design studies whenever applicable.

Part IV

CONCLUSION AND PERSPECTIVES

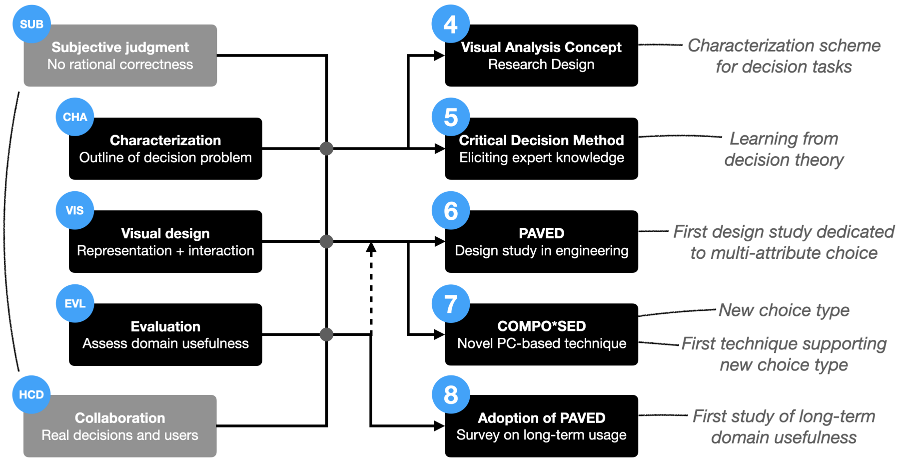

Figure 9.1: Research challenges addressed in this thesis (left). The answers to the challenges C_{CHA}, C_{VIS}, and C_{EVL} are provided in the respective chapters (right). The overarching challenges C_{SUB} and C_{HCD} generally influenced the concepts and techniques presented.

9

CONCLUSION

THIS chapter summarizes the findings, limitations, and contributions of this dissertation.

9.1 SUMMARY OF FINDINGS

Part II provided foundations from human science, real-world applications, and information visualization.

Chapter 2 (Multi-Attribute Choice) reviewed the discourse on decision-making in human science as well as engineering design as an exemplary real-world application. We found that human scientists agree on the basic ingredients of a decision task. The view on how people make decisions, however, has shifted: from normative decision theory, which attributes ideal rational-economic behavior to decision-makers, to descriptive decision theory, which attributes subjective and potentially irrational behavior to decision-makers. We classify multi-attribute choice as a constructive problem, where decision-makers make judgments on the fly rather than drawing from existing well-defined preferences. While both user groups engage in data exploration, decision-makers differ from analysts in that they collect evidence for their decision rather than highlight (any) interesting trends and patterns. Finally, we realized that engineering design as an application domain shares a lot with constructive decision models. With interacting components in systems, the domain even introduces a new type of decision problem. We also found a significant number of interactive visualizations for expert decision support to originate from the field of engineering design.

Chapter 3 (State of the Art) carved out the research gaps associated with visually assisting multi-attribute choice by contrasting the identified research challenges with existing works in information visualization. Our review of design study methodologies revealed a lack of guidance on how to systematically elicit expert knowledge and strategies involved with decision-making and how to validate the (long-term) usefulness of a visualization tool in real-world decision settings. We found multivariate visualizations without dedication to decision support to be effective for basic analytic activities but to fall short on trade-off analyses needed for decision-making. Visualizations meant for general-purpose decision support address decision-specific requirements but might not meet expert decision-makers' needs for operational and visual flexibility. Visualization design studies on expert decisions typically lack a dedicated consideration of multi-attribute choice tasks. Some rely on disadvantageous dimension reduction or weighting approaches, while others do not depict complete alternatives or find a desirable region rather than a single option. Most do not support advanced decision-specific activities like incrementally constructing preferences and reconciling conflicting information to make a final choice. No design study targets decisions where choices affect each other. Finally, interaction mechanisms provide effective support of visualization-based decisions but could benefit from integrating specific knowledge about the decision problem.

Part **III** built upon the theoretical and empirical foundations summarized in Part II to provide a concept for visual analysis for multi-attribute choice and demonstrate its application to different visualization challenges raised by choice problems (Figure 9.1).

Chapter 4 (Concept) presented the research goals, targeted decision problem, and research design of this thesis by considering previous findings from data science, human science, and information visualization. Properties of multi-attribute choice regarding data, users, and tasks were structured in a characterization scheme to untangle the diversity of finding the best among a finite set of multi-attribute options in personal and professional contexts. On this basis, the targeted decision problem for this thesis has been defined. Studying such problems in real-world settings calls for a qualitative research approach. It connects the visualization researcher's design process to the decision-maker's constructive decision process. As orthogonal dimensions, both processes span our conceptual space, which also characterizes the range of contributions to visual analysis for multi-attribute choice.

Chapter 5 (Critical Decision Method) addressed the lack of methodological prescriptions for eliciting domain knowledge in real-world situations (as identified in Section 3.1.2) by investigating an interview technique from cognitive science. The five interview steps of the *Critical Decision Method* provide a systematic way of learning about the conditions framing multi-attribute choice performance. This opens up a novel perspective on characterizing multi-attribute choices by favoring decision points over tasks. The focus on real-world incidents aligns well with the required realism in understanding work practices. Still subject to validation, the technique also makes it easier for decision-makers to reflect on and narrate their cognitive processes.

Chapter 6 (PAVED) addressed the lack of visualization solutions for real-world multi-attribute choices (as identified in Section 3.2.3) by presenting a visualization design study in the field of engineering design. It confirmed previous findings that parallel coordinates are advantageous when introduced in a new application area. A careful domain characterization resulted in a decision process model fed by multiple visualization task taxonomies. The observed decision process followed an elimination-by-aspects strategy. We concluded that the visual design should prefer simple over flexible interaction and a lossless projection over dimension reduction. These rationales informed the development of *PAVED*, an interactive parallel coordinates visualization that supports multi-attribute choices with hundreds of options and tens of attributes. For simplified interaction where the optimization direction is known, we proposed preference brushes, which are tied to the high-quality end of an axis. As confirmed by a qualitative field study, the visualization helped engineers learn what level of performance is achievable under different conditions. Reflections on the design process revealed that the domain experts valued overview and transparency more than guidance and that visualization researchers should take the experts' suggestions seriously, even if they feel that their visualization expertise is underrated. Although designed for professional decision-makers, we assume the visualization to also work for casual decision-makers, e.g., consumers, due to its intended simplicity. Finally, the design study revealed that visualization support is needed for making decisions where choices affect each other.

Chapter 7 (COMPO*SED) addressed the lack of visualization support for co-dependent choices (as identified in Chapter 6) by introducing a novel visualization technique. Conventional parallel coordinates cannot serve the dual role of attributes that represent individual subject properties and contribute to interoperability or overall performance at the same time. To address this challenge, the chapter presented *Composite Parallel Coordinates*. They help decision-makers work through a series of co-dependent multi-attribute choices to choose the combination of variants that optimizes the overall performance. Variants of single components are depicted in juxtaposed parallel coordinates plots. A vertical layout has shown to be superior to horizontal or circular layout. Where design parameters or criteria refer to multiple components, the associated axes can be superimposed. The option selectors are extended by dedicated interaction patterns, such as locks, to allow for simultaneous exploration of unique component properties and emergent system properties. Two usage scenarios and one case study illustrated how *COMPO*SED* can be used to make co-dependent choices. Unlike common practice, we performed the case study with a pair of experts. Watching the experts learn not only from the data set but also from each other provided insights that we would not have gained otherwise.

Chapter 8 (Long-Term Adoption) addressed the lack of works observing the long-term domain usefulness of a decision support tool in terms of adoption and routine use. Four years after having completed our design study on Pareto front visualization (Chapter 6), we wanted to know whether the resulting tool is still used in the experts' daily workflows without our involvement. In our online survey, four out of ten target users reported to have used PAVED routinely and on their own initiative to make or demonstrate choices about motor designs. The adopters' qualitative feedback and usability ratings confirmed our design rationales regarding a lossless overview of all options and reduced yet effective interaction mechanisms. The main reasons for unsuccessful adoption were missing awareness or need of the tool. Only one participant actively refused the use due to an integration issue. Despite the limited group of domain experts in our collaboration, their repeated, self-initiated use adds evidence that our visualization in fact addresses their true needs. It also supports our design approach of preferring simple visual designs to address a domain problem over complex encodings for novelty's sake.

9.2 DISCUSSION OF LIMITATIONS

The purpose of this section is to briefly comment on the general limitations of the work presented in this thesis that go beyond the more particular limitations discussed in the respective chapters above.

Objectivity We do not view subjectivity as a general threat to validity. Still, our practice carries the risk of subjective bias in certain situations. One example is the survey of background works in information visualization (Chapter 3). Our literature review was informed by recent visualization surveys on decision-making, readings, and discussions with fellow researchers. The final selection of relevant works was based on the subjective judgment of a single researcher (the author). Without a clear understanding of decision tasks among visualization researchers, there are many ways to synthesize the selected works. We largely built upon structures applied in previous surveys. While this favors objectivity and coherence across literature reviews, we might have overlooked newly emerging topics. Another example is the characterization scheme for multi-attribute choice (Chapter 4).

Its main limitation is that it is not grounded in a systematic literature review but evolved from our experience gained throughout the years. As such, it might involve a subjective bias. Further examples from literature and an independent thematic coding of the properties are needed to validate the characterization scheme. In the remaining chapters (Chapters 5 to 8), subjectivity has been addressed by reflecting on the findings with respect to the context of their construction. Despite these aspects, we agree that personal involvement of the researcher is central and desired in design study methodology [243].

Interdisciplinarity Visualization research offers technologies that can be applied to problems in a different (scientific) domain. In the context of this thesis, visualization research intersects with the topic of decision-making. Combining the specialized knowledge from both subjects offers considerable potential to solve complex real-world problems. Learning from a discipline that has devoted a long history to studying decision-making prevents visualization research from re-inventing the wheel [67]. Still, the ability to assess the relevance and implications of concepts from a foreign discipline requires visualization researchers to understand the way of thinking and the terminology that are common in that discipline. Despite our effort to identify and present findings on how humans make decisions in an accurate manner, a cross-disciplinary picture of the collective evidence in decision-making is out of the scope of this thesis. In choosing appropriate decision-related concepts for our visualization research, the significance of models and methods was sometimes difficult to judge without a profound background in psychology or economics. Transferring concepts from decision theory to our visualization applications generally carries the risk of loosing nuances of the original understanding. Still, we embrace our interdisciplinary research approach as a valuable source of knowledge.

Generalizability Applied visualization research is receiving increased attention. The strength of this thesis is its major limitation at the same time: our findings have been generated using a qualitative research approach on real users, data, and tasks in an application domain. Lessons learned are only reliable in this specific context and need to be treated as assumptions beyond that. In contrast to quantitative research, where testing hypotheses in experiments produces inherently generalizable results, qualitative research is about developing an in-depth understanding of a topic based on interviews, observations, and literature reviews. Investigating how information visualization can assist decision tasks on a case-by-case basis produces rich yet potentially isolated insights. Moving from individual observations to a contribution that provides value to the visualization community can be achieved in different ways [168, 236]. In general, abstraction, reflection, and a synthesis of results with existing knowledge help obtain results that transfer beyond individual design studies.

Novelty As is the case with many design studies, our technical contributions mainly resulted from a meaningful combination of existing building blocks into solutions that are tailored to real-world decision problems. We explained in our conceptual approach (Chapter 4) that, while we consider the novelty and innovative character of an approach important, our focus is on leveraging existing visualization and interaction techniques to offer effective solutions for given real-world decisions. This involves careful problem abstractions including typical data characteristics and tasks that also advance the general understanding of decision-making as a sub-field of visualization. As such, our works can significantly contribute to visualization research. For example, they can adapt proven methods from other disciplines to help visualization design practice as in the case of the Critical Decision Method (Chapter 5). They might also solve a more general problem underlying the domain-specific task as in the case of PAVED (Chapter 6). Finally, they might solve a problem that is relevant but has not been addressed before as in the case of COMPO*SED (Chapter 7).

9.3 THESIS CONCLUSION

Decision-making is an ambiguous subfield of visualization. On the one hand, researchers have named it a core goal of visualization from the earliest days. On the other hand, despite their consensus on its importance, decision-making is underspecified in the sense that references in visualization theory and practice often remain at an introductory or overview level. Although some researchers have been working on an explicit tie to visualization research, there is a lack of experience how visualization and decision-making activities interact in real-world situations.

The research presented in this thesis provides answers regarding how interactive visualizations can effectively assist experts in making real-world choices among multi-attribute alternatives. Previous works rarely dedicated field work to study this question. We found application-oriented visualization research to generally lack an emphasis on decision processes. To approach this gap, our research design provides for a joint consideration of the visualization researcher's design process and the decision-maker's decision process. For this purpose, we built upon a close collaboration with real decision-makers.

Previous definitions and characterizations of decision tasks for visualization were not grounded in real-world settings (**RQ_1**). As a result, they were barely operationalizable. Every decision is as unique as the person who makes it and a universal definition is out of our reach. Still, this thesis provides precise characterizations and abstractions of expert choices in different domains that draw a picture of the information needs of real decision-makers. Connecting our experience in the wild to a review of models from decision theory, we have extended the existing multi-attribute choice definition by classifying it as a constructive problem, where the preferences of decision-makers unfold throughout their decision process. Despite an understanding of the task to be designed for, the lack of a clear, decision-focused terminology made it difficult to precisely claim the value of visualization for decision support. To make it easier for visualization researchers to define their design target, we have distilled a compact vocabulary comprising pairs of data-user-task properties that can be used to describe a targeted decision problem. Finally, we proposed a formalization of co-dependent choices as a novel type of decision that had not been defined before.

While visualization research has embraced the approach of characterizing a domain problem to inform visual encoding decisions, problem characterizations lacked a dedicated consideration of decision-making activities (**RQ2**). In particular, methods to access the cognition, domain knowledge, and experience of expert decision-makers were rare. The domain characterization method proposed in this dissertation offers a decision-oriented approach by adopting an interview technique originally developed to elicit expert knowledge in naturalistic decision-making. Its procedure does not only provide concrete steps to capture the knowledge and experience underlying real-world decision-making. It also offers general prescriptive guidance that domain characterizations in design study methodology have lacked so far. By investigating the conditions of the technique's applicability in visualization research, we further contribute to connecting the disciplines of decision theory and information visualization.

Given the task of selecting the most preferred option among multi-attribute alternatives, not many decision tools (**RQ3**) and their evaluations (**RQ4**) have found their way into real-world application. On the basis of our rationales around visual analysis for multi-attribute choice, we introduced two interactive parallel coordinates visualizations serving as decision support tools for real-world expert decisions. They resulted from applying our concepts in two design study projects that we conducted on increasingly complex choices in the field of engineering design. While the use of parallel coordinates is not novel for decision support, PAVED provides a simple yet effective means for engineers to construct and apply preferences as they learn what level of performance is achievable under which conditions. Its extension COMPO*SED is the first tool to help decision-makers explore the side effects of co-dependent decisions. Its design principles generally apply to the visual analysis of increasingly important multi-model data, e.g., originating from simulating interacting phenomena. Both tools have been evaluated with domain experts on their day-to-day professional decision problems, providing useful evidence regarding the validity of visualization idioms within and beyond particular decision contexts. Besides their contributions to the body of visualizations that advance

Table 9.1: By comparing the research challenges (Sec. 1.2) to prior works, we identified research gaps (Sec. 3.3), from which we derived research goals (Sec. 4.2) that informed our concept and contributions.

Challenges	**Research Gaps**	**Goals**	**Contributions**
Subjective judgment	Lack of consideration of constructive preferences	Design for constructive preferences	Overarching
Characterization of decision-making activities	Lack of systematic elicitation of tacit knowledge and decision strategies, lack of explicit distinction between decision and analytic tasks	Characterize multi-attribute choice, learn from other disciplines how to study decision tasks	Characterization scheme (Ch. 4), knowledge elicitation method (Ch. 5)
Visual design and interaction	Lack of visualization designs tailored to constructive preferences, lack of designs for decision problems where trade-offs affect each other	Tailor visualization designs to constructive preferences, support co-dependent choices	Decision tools PAVED (Ch. 6) and COMPO*SED (Ch. 7)
Evaluation of decision support	Lack of validation on expert decisions in real-world settings, lack of observation of long-term usefulness	Validate visualization designs in the wild	Short-term field studies (Ch. 6 and 7), adoption study (Ch. 8)
Collaboration with decision-makers	Lack of design studies on multi-attribute choice	Apply user-centered design	Overarching

decision-making activities on large data sets, e.g., constructing preferences or reconciling conflicting information, the tools demonstrate the general applicability of our concepts for multi-attribute choice.

Table 9.1 summarizes the presented contributions in response to our research goals that emerged from the identified challenges and gaps. By conducting the research presented in this thesis, we have learned

- what real-world decisions look like and how we can describe them for the purpose of visualization design,
- how we can elicit the cognition, domain knowledge, and experience underlying these decisions,
- what is important when collaborating with domain experts and designing visualization tools for decision support, and
- how to assess their usefulness under real-world conditions.

9.4 SUMMARY OF BENEFIT

We anticipate that a wide range of readers who are involved with decision-making or visualization will find this dissertation useful. Researchers in decision-related disciplines might get a sense of how their scientific concepts impact research in other fields. It might also provide them with a new perspective on their own research, potentially helping them break new ground in their field. Visualization researchers addressing decision problems will find groundwork on assisting real-world decisions with visualizations. Our lessons learned provide guidance regarding how user-centered design can inform the development of meaningful visualization tools, with a focus on but not limited to decision problems. Visualization research might generally benefit from increased awareness of decision tasks, potentially leading to more rigorous decision support claims. Decision-makers themselves, whether they are facing scientific, industrial, or personal decisions, will find solutions to decision tasks they might know in this or similar forms. They might also develop an idea of how application-oriented visualization research can help solve their problems, effectively narrowing the knowledge gap between visualization researchers and target users. Finally, we anticipate that this thesis might also serve as a general source of inspiration for other PhD candidates.

10 OUTLOOK

Decision support is a traditional and at the same time emerging topic in visualization. In this thesis, we conducted research on concepts, methodologies, and techniques to assist expert decision-makers with interactive visualizations. We believe that academic research initiates as many questions as it answers. This chapter highlights opportunities for further practice that emerge from the presented contributions. Future research directions regarding the individual methods and techniques have been discussed in Chapters 5 to 8.

The characterization scheme for multi-attribute choice tasks that we introduced in Section 4.3 provides a first starting point for exploring future work directions. New opportunities and challenges arise from addressing varied properties of the targeted choice task (Table 4.1). For example, researchers could aim at supporting groups instead of individual decision-makers, exploring progressively generated rather than previously known alternatives, or dealing with decisions where outcomes might be uncertain rather than certain. Besides targeting variations of the choice task, further research could also provide additional contributions throughout the range spanned by the design study and decision process axes (Figure 4.6). Based on our learnings from this thesis, we provide selected recommendations in this regard.

Adopt more models and methods from decision theory In this thesis, we have connected to findings from decision theory in two ways. First, we have learned from applying the Critical Decision Method that favoring decision points over analysis tasks opens up a novel perspective on decision-focused problem characterization (Chapter 5). Second, by classifying the decision strategy of our domain expert based on the framework proposed by Payne et al. (see Section 2.1.4), we realized the relevance of filtering for choices in engineering design and chose our main design rationale accordingly (Chapter 6). We believe that visualization research will benefit from investing further efforts in investigating whether (and which) existing models and methods in decision theory are indeed appropriate for serving the data analysis and visualization needs present in our field [67]. One way to achieve such appropriation is to collect further evidence by reflecting on the application of carefully selected Cognitive Task Analysis methods [178] to elicit domain knowledge for the purpose of informing visualization designs for various data analysis settings. Another way is to further investigate the role and potential of known decision strategies (Section 2.1.4) in visualization research. Visualization researchers could classify their domain experts' approaches according to these strategies as part of their problem characterization. This also facilitates future meta studies that might investigate what decision strategies are supported by current visualization tools for decision-making. Torsney-Weir et al. proposed an exemplary meta study on 21 visual parameter space analysis approaches [269]. Synthesizing the findings of these efforts could ultimately contribute to the integration of decision tasks into visualization theory, e.g., task taxonomies, to improve the information base for rigorous visualization design decisions.

Engage in adaptive decision support This future work direction partly relates to the above. Decision theory has provided evidence that people switch among a variety of decision strategies in response to varying task context and complexity [201, 203]. Adaptive visualizations dynamically change the visual presentation of data in response to varying characteristics or needs collected about the user [3]. This information can be explicitly provided (by the user) or inferred from the trace of user actions. We consider decision-makers switching their strategy during a choice task as a symptom of changed information needs. Profiling the decision-makers personality might additionally allow for conclusions about how to interpret these symptoms. We hypothesize that meeting the varying needs by adaptive visualization holds the potential to improve the overall choice task performance. These thoughts have several implications for visualization research. First, problem characterization needs to consider decision-making as an adaptive process rather than a linear (even if iterative) sequence of low-level tasks. Again, findings from decision theory can help get started with this. Second, visual designs need to carefully integrate the adaptation needs carved out in the problem characterization with (explicit or implicit) user feedback during analysis to adapt visual representations and interactions accordingly. Adaptation targets could be, for example, visual encodings, navigation strategies, levels of detail, or color mappings. Provenance approaches can help keep track of user actions and gained insights on which a decision is based [300]. Third, evaluations need to ensure that unexpected results do not stem from participants changing strategies differently from what researchers expected based on the problem characterization.

Develop dual-use decision support Visualizations typically serve one of three major goals: exploratory analysis, confirmatory analysis, or presentation [235]. Like most visualization research related to decision-making, our work investigated problems in the context of an exploratory analysis. However, from our design study on PAVED (Chapter 6), we learned about the importance of being able to transparently communicate how one arrived at a certain decision. This holds especially true if decisions are made on behalf of others who, e.g., do not have the required expertise. Whether these parties approve the decision or wish to refine it is likely associated with uncertainty. Preparing the essential information to replay and revisit a decision process together with an affected stakeholder makes the visualization goal turn from exploratory analysis into presentation. Depending on the involvement of the stakeholder, a presentation might also carry a notion of group decision-making. For a seamless transition, future research might investigate visual representations and interactions as "dual-use" techniques, in the sense of techniques that can be used in the different contexts of exploratory analysis and presentation. Depending on the prior knowledge of the stakeholder and the level of explanation by the decision-maker, simplified representation and functionality is likely needed for presentation. For example, simplification could mean to reduce the number of depicted alternatives, specify certain presets or filters that can otherwise be interactively modified, or freeze at a certain level of detail where further in-depth comparison would otherwise be possible. In this sense, dual-use visualizations could also be considered a form of adaptive decision support, where visualizations adapt to the changing needs of decision-makers and stakeholders as the task changes from exploratory analysis to presentation.

Develop decision quality metrics for evaluation From the application of qualitative evaluation methods in the context of our design study perspective, we have collected rich and meaningful insights into how experts approach decision-making and how they experience the use of a visualization to assist their day-to-day choice tasks. Still, accurately evaluating visualizations for their ability to support choice tasks is generally difficult because these tasks typically have no clear ground truth. Objective metrics for decision quality are rare. Dimara et al. proposed an accuracy metric that captures the consistency between the choice made and self-reported preferences regarding the importance and optimization direction of attributes [64]. However, they found no clear difference in decision accuracy between three studied visualization techniques (parallel coordinates, scatterplot matrix, and tabular visualization). They hypothesized that the accuracy metric might suffer from noise, which causes it to overlook subtle but meaningful differences between the techniques. For example, participants might not have been able to perfectly express their preferences by specifying attribute importance and optimization direction only. It would therefore be valuable to investigate how the sensitivity of decision accuracy metrics can be increased. The experiment could be repeated with a preference elicitation method that is assumed to be more accurate. In the context of human-centered ranking creation, we studied the application of attribute scoring functions [233], which offer a more fine-grained specification of preferences on attribute values [234]. Compared to the original preference elicitation using optimization direction, we expect findings of the repeated experiment to hint at an increased ability of the decision accuracy metric to capture subtle differences between the visualization techniques. Such a sensitive objective decision accuracy metric would greatly complement the qualitative feedback in a mixed methods approach to validating visualizations for their ability to assist choices in the wild. Once we know how to accurately measure decision quality, we could also study its relation with the diversity of exploration activities, in which a decision-maker engages during the decision process. People might make more accurate and trustworthy decisions when they develop a differentiated understanding by taking different perspectives, i.e., engaging in numerous and diverse exploration activities, rather than quickly jumping to conclusions. A first step towards such studies has been made by identifying behavioral variables for ranking tasks [21].

Synthesize findings from future design studies This final research direction is a call to holistically reflect on collaborations with expert decision-makers to improve visualization guidelines. We successfully applied our complete design study approach to two increasingly complex decision problems in the field of engineering design, each involving a careful problem characterization and a (long-term) validated visualization design. As indicated above, other users, data, and decision tasks yield additional promising configurations for future visualization support. Beyond engineering design, a variety of other application fields can benefit from the solutions presented, like consumer research, finance, or automotive supply. Other visualization researchers can benefit from our work in applying the design study methodology to develop useful visualizations that assist expert decision-makers from such domains in choosing among large sets of multi-attribute alternatives. Since decision tasks have not received much dedicated attention in visualization theory and research, it will be interesting to investigate the generalizability and adaptability of our approach. Similar to the meta studies of general design studies [243] or design studies on parameter space analysis [240], further design studies on decision-making yield the potential to reflect on and synthesize the experience and lessons learned across use cases to move towards more systematic and rigorous designs and evaluations of decision-focused visualizations.

Part V

APPENDIX

A RESULTS OF THE CRITICAL DECISION METHOD

A.1 ELECTRIC DRIVE DESIGN

Table A.1: Situation assessment record (SAR) for electric drive design.

Situation Assessment 1	Plausibility check
Cues	Invalid solutions, deviation from specifications
Experience	Acceptable proportion of invalid/valid solutions
Goals	Ensure valid solutions and achievable targets
Decision Point 1	If targets out of reach, adjust input parameters
Situation Assessment 2	Adjust input parameters (elaboration)
Cues	Input parameter ranges of undesired solutions
Goals	Generate more valid solutions in the next run
Decision Point 2	Refine input parameter setting
Situation Assessment 3	Convergence check (shift)
Cues	Hypervolume stable over several generations
Goals	Determine whether evolution has come to a halt
Decision Point 3	If optimization converged, proceed
Situation Assessment 4	Satisfaction of hard constraints (shift)
Cues	Well-defined limits specified by customer
Knowledge	Customers typically want good efficiency, i.e. low losses, magnet mass and price correlate
Experience	Expectable power loss under different conditions
Goals	Meet non-negotiable customer requirements
Decision Point 4	Exclude individuals not meeting requirements
Situation Assessment 5	Further trade-offs and final choice (elaboration)
Cues	Important criteria, criteria correlations, individuals becoming (de-)selected upon filtering
Basis	Apply well-defined constraints followed by underspecified constraints, observe effects of filters
Knowledge	Operational characteristics suggest to restrict harmonic distortion and cogging torque
Experience	Underspecified criteria, desired trade-offs from technical perspective, flexibility of specifications
Goals	Optimize performance within customer specifications
Decision Point 5	Proceed with most preferred drive design

Table A.2: Decision requirements table for electric drive design. The rows correspond to the decision points identified in the SAR (Table A.1).

	What is the decision?	**Why is it difficult?**	**How is it made?**	**What is the aid? How does it help?**
1	Determine whether computed solutions are valid or satisfactory	Vague specification of when the customer specification is in reach, acceptable proportion of valid and invalid individuals not obvious	Compare the targets specified by the customer with the target ranges covered by the optimization, observe the proportion of valid and invalid individuals across generations	Scatter plot axes and distribution of individuals in scatter plots convey target ranges, a plot depicts the proportion of valid and invalid individuals in each generation
2	Adjust input parameters	No inverse function $Y \rightarrow X$ that tells us how to choose the input parameters to achieve a desired outcome	Look at input parameter values associated with undesired target values	Histograms highlight input parameter ranges that are associated with undesired solutions
3	Determine convergence of optimization	Difficult to understand whether a genetic algorithm has converged to an optimum, especially with multiple criteria	Observe whether the hypervolume remains stable over multiple generations	A line chart depicting the development of a hypervolume metric
4	Satisfy customer specifications	This is rather easy, because it is about applying well-defined constraints	Use the specifications from the customer order and obvious engineering indications as filters	Scatterplot matrix conveying (remaining) distribution of individuals, experience about which level of efficiency is appropriate
5	Choose most preferred drive design	Improving on one criterion will worsen another criterion	Explore different what-if scenarios and compare their outcomes	Interactive filtering in scatter plots

A.2 THERMAL POWER PLANT OPERATION

Table A.3: Situation assessment record (SAR) for power plant operation.

Situation Assessment 1	Plausibility check
Cues	Deviation from on-site measurements
Goals	Validate simulation model
Decision Point 1	If model is valid, proceed
Situation Assessment 2	Exploration of reference case (shift)
Cues	Distributions of velocity, temperature, oxygen, NOx, and CO across boiler, co-occurrence of high oxygen and high temperature, peaks in CO concentration
Knowledge	NOx forms at high oxygen and temperature, oxygen distribution influences NOx distribution, CO indicates where fuel does not burn correctly
Experience	Most interesting indicators are velocity, temperature, and oxygen, simulation carries uncertainty
Goals	Understand strengths and weaknesses and how to modify operation mode to achieve targets
Decision Point 2	If optimization potential is understood, proceed
Situation Assessment 3	Optimization of reference case (elaboration)
Cues	Temperature (oxygen, NOx) distributions for reference case and operation modes, total CO and NOx emissions
Knowledge	Hard constraints: NOx emissions and exit temperature should remain the same
Experience	Insights from SA 2, variation of local temperatures can be tolerated as long as exit temperature is met, trade slightly increased NOx emissions for significant decrease of unburned fuel
Goals	Determine most preferred operation mode: reduce unburned fuel, maintain transferred heat, NOx emissions, and exit temperature
Decision Point 3	Collect preferred operation modes
Situation Assessment 4	Final decision and confirmation (elaboration)
Cues	Local optimum, robust to small changes of operational parameters, approaching deadline, budget
Goals	Maintain reasonable cost-benefit ratio, provide recommendations for on-site adjustments
Decision Point 4	Present two selected configurations to customer

Table A.4: Decision requirements table for power plant operation. The rows correspond to the decision points identified in the SAR (Table A.3).

	What is the decision?	Why is it difficult?	How is it made?	What is the aid? How does it help?
1	Determine whether simulation results are plausible	Not obvious what simulation error is acceptable	Compare real-world measurements of reference case with simulated values	Table with side-by-side numbers facilitates comparison
2	Determine optimization potential	Complex operation behavior, multiple influencing factors need to be considered	Simulate reference case and investigate its performance	Side-by-side comparison of oxygen, temperature and NOx heatmaps helps find co-occurrence of high oxygen and high temperature producing NOx emissions, CO heatmap helps see where fuel does not burn properly
3	Determine preferred operational parameter settings	Fluid problems based on non-linear equations and highly coupled effects, trying to fix one problem might introduce another problem, it is difficult to find an equilibrium between all the involved effects	Comparison of temperature (oxygen, NOx) heatmaps for reference case and different configurations of the operation mode	Side-by-side arrangement of heatmaps/numbers supports direct comparison of available options
4	Confirm decision and finish the optimization	Difficult to determine whether a solution is good enough, potential of further improvement largely based on intuition	Sensitivity analysis: if small changes produce similar results, the solution is the most preferred one, consider approaching deadlines or budget limitations	n.a.

B

TASK ABSTRACTION FOR THE DESIGN OF PAVED

We provide additional information on the task abstraction of the design study on PAVED (Chapter 6). We detail the process of identifying relevant visualization tasks and provide a comprehensive list of possible analysis questions related to each task.

B.1 METHODOLOGY

We considered two sources of information: 1) elicited domain-specific decision tasks involved with electric drive design and 2) existing task taxonomies that have been proposed in visualization literature.

We performed a literature review of 11 papers published between 1990 and 2013 (Table B.1) that provide different perspectives on visualization tasks. From these works, we collected the proposed high-level and low-level visualization tasks, initially without paying attention to duplicates or their applicability to our topic. We then excluded duplicates (including synonyms), resulting in a collection of 54 unique visualization tasks. These were further filtered according to their relevance for our domain characterization, resulting in a final collection of 10 tasks. We mapped these tasks to the domain-specific decision process by arranging them along two orthogonal axes: abstraction level and stage in the decision process. The resulting task abstraction is shown in Figure 6.3.

Table B.1: 11 taxonomies serving as a baseline for the task abstraction.

Year	**Authors**	**Abbreviation**	**Ref.**
2013	Brehmer and Munzner	Topology of Abstract Tasks	[38]
2013	Schulz et al.	Design Space of Tasks	[235]
2009	Munzner	Nested Model	[185]
2006	Valiati et al.	Taxonomy of Tasks	[283]
2005	Card and Pirolli	Sensemaking Process	[210]
2005	Amar and Stasko	Knowledge Precepts	[7]
2005	Amar et al.	Low-Level Components	[8]
1999	Card et al.	Knowledge Crystallization	[42]
1998	Zhou and Feiner	Visual Accomplishments	[311]
1996	Shneiderman	Task by Data Type Taxonomy	[245]
1990	Wehrend and Lewis	User Cognitive Tasks	[292]
1990	Roth and Mattis	Display Functions	[227]

B.2 SELECTED TASKS AND UNDERLYING ANALYSIS QUESTIONS

T1 Inform

- Validation of the simulation
  - Does the simulation model produce plausible results?
- Overview of the solution space
  - How diverse are the solutions?
  - Are there clusters of similar solutions?
  - Are there outliers?
- Overview of the objective space
  - What is the distribution of solutions within each objective range?
  - What is the nature of conflicts between objectives?

T2 Identify > T2.1 Search > T2.1.1 Browse

- Region of interest
  - What are combinations of desired objective ranges (e.g., that include a known successful solution)?
  - What types of trade-offs between objective values are available?
  - What is the ranking of solutions for each objective?
  - Which topology provides the majority of Pareto-optimal solutions?
  - How do the objective values change when moving through the solution space/region of interest?
- Navigation based on a focal solution
  - What are the absolute values of a solution?
  - Which alternative can be selected should a gain in a certain objective be desired?
  - What are options of gain and with what losses do they come?
  - How do other solutions differ from the focal solution?
- Navigation through the preferences space
  - What is the effect of preference variations on the best choice of the solution?
  - For how long does a solution stay in the region of interest?
  - For how long does a topology stay the most preferable?

T2 Identify > T2.1 Search > T2.1.2 Relate

- With respect to objectives
  - How are good or bad values of objective X related to good or bad values of the remaining objectives?
  - How do changes in objective X affect the remaining objectives?
  - How much can I gain in objective X when accepting worsening in objective Y?

- With respect to solutions
  - On what (combinations of) objectives is a solution superior/dominated?
  - What is the relation between two subsets of solutions?
  - Where is the focal solution located within a subset of solutions (e.g., with respect to each objective)?
  - How are design parameters related to a subset of solutions of interest?
- With respect to topologies
  - What is the relation between two topologies?
  - Which topologies perform well for which objectives?

T2 Identify > T2.1 Search > T2.1.3 Filter

- Objectives of interest
  - Which objectives can be excluded to reduce complexity?
  - Which pairs of objectives exhibit the most critical conflicts?
- Region of interest according to trade-offs
  - Which solutions implement known or desired trade-offs?
  - Which solutions are superior w.r.t. a global quality score?
- Region of interest according to preferences
  - Which solutions lie within desired value ranges?
  - Which solutions satisfy the joint preferences of different stakeholders?
- Undesired regions of the solution space
  - Which solutions are considered undesired (e.g., infeasible, extremal values, or too expensive) and can be eliminated?

T2 Identify > T2.2 Compare

- Compare solutions based on their objective values
  - What is the trade-off between two solutions?
  - Which solution has better values on most objectives?
  - Which improvements/deteriorations in individual objectives cancel each other out?
  - How are solutions different w.r.t. metrics like sensitivity?
- Compare similar solutions or different sets of solutions
  - How do solutions that are similar w.r.t. a subset of objectives perform in the remaining objectives?
  - What is the performance difference between two topologies?
  - How are similar solutions different in their sizes of the dominated areas in the objective space?

T3 Confirm

- Quality
  - Does the selected solution satisfy the desired constraints in the most important objectives?

  - What are the solution's values for objectives that were not considered for the decision?
  - How well does the selected solution perform compared to other favorite solutions?
- Sensitivity
  - How does the quality of a solution change as design parameters are slightly varied?
- Confidence
  - What is the confidence in the selected solution?
  - What additional analyses (e.g., simulations) should be performed?
  - What consequences (e.g., manufacturing costs, etc.) might arise from choosing the selected solution?
  - What will the physical instantiation of a solution look like?

T4 Verify + T5 Communicate

- Decision
  - How did the decision evolve?
  - On what insights is the decision based?
  - What alternative solutions to the one provided were encountered during the search?
  - Why is the suggested solution superior to the alternatives?
- Decision refinement
  - What alternatives are reasonable w.r.t. refined preferences of other stakeholders after having observed the decision made by the decision-maker?
  - What alternatives fulfill additional qualitative conditions (e.g., required post-processing, the accessibility of a surrounding environment, etc.)?
  - How does a reduction in budget affect the achievable performance?
- General recommendations
  - Why should certain solutions be avoided?
  - What is the decision-making scope given required hard constraints and an available budget?
  - Which topology can be used for which type of intended use?
  - Which budget should be made available at least to reach a satisfying performance?
  - What are utopian visions of what is possible?

C RESEARCH EXPERIENCE

C.1 KEY PUBLICATIONS

- **L. Cibulski**, T. May · Revisiting PAVED: Studying Tool Adoption After Four Years · *EuroVis Short Papers* · 2024
- C. M. Barth, J. Schmid, I. Al Hazwani, M. Sachdeva, **L. Cibulski**, J. Bernard · How Applicable are Attribute-Based Approaches for Human-Centered Ranking Creation? · *Computers & Graphics, 114* · 2022
- **L. Cibulski**, E. Dimara, S. Hermawati, J. Kohlhammer · Supporting Domain Characterization in Visualization Design Studies With the Critical Decision Method · *IEEE VisGuides: Visualization Guidelines in Research, Design, and Education* · 2022
- **L. Cibulski**, T. May, J. Schmidt, J. Kohlhammer · COMPO*SED: Composite Parallel Coordinates for Co-Dependent Multi-Attribute Choices · *IEEE Transactions on Visualization and Computer Graphics, 29(10)* · 2022
- **L. Cibulski**, J. Schmidt, W. Aigner · Reflections on Visualization Research Projects in the Manufacturing Industry · *IEEE Computer Graphics and Applications, 42(2)* · 2022
- J. Schmid, **L. Cibulski**, I. Al Hawazmi, J. Bernard · RankASco: A Visual Analytics Approach to Leverage Attribute-Based User Preferences for Item Rankings · *EuroVis Workshop on Visual Analytics* · 2022
- S. Hermawati, **L. Cibulski**, G. Lawson · Towards Using the Critical Decision Method for Studying Visualization-Based Decision-Making · *Ergonomics and Human Factors* · 2021
- **L. Cibulski**, H. Mitterhofer, T. May, J. Kohlhammer · PAVED: Pareto Front Visualization for Engineering Design · *Computer Graphics Forum, 39(3)* · 2020

C.2 FURTHER PUBLICATIONS

- P. Louis, **L. Cibulski**, J. Suschnigg, E. Marth, H. Mitterhofer, J. Kohlhammer, T. Schreck, B. Mutlu · Visual Exploration and Analysis of Simulation and Testing Data in Motor Engineering · *IEEE Computer Graphics and Applications* · 2024

- **L. Cibulski**, H. Mitterhofer · Interactive Visualization: A Clear View of Product Design · *WT Werkstattstechnik, 111(4)* · 2021
- **L. Cibulski**, T. May, B. Preim, J. Bernard, J. Kohlhammer · Visualizing Time Series Consistency for Feature Selection · *Journal of WSCG, 27(2)* · 2019
- S. Alemzadeh, T. Hielscher, U. Niemann, **L. Cibulski**, T. Ittermann, H. Völzke, M. Spiliopoulou, B. Preim · Subpopulation Discovery and Validation in Epidemiological Data · *EuroVis Workshop on Visual Analytics* · 2017
- **L. Cibulski**, B. Klarin, M. Sopouch, B. Preim, H. Theisel, K. Matković · Super-Ensembler: Interactive Visual Analysis of Data Surface Sets · *Spring Conference on Computer Graphics* · 2017
- **L. Cibulski**, D. Gračanin, A. Diehl, R. Splechtna, M. Elshehaly, C. Delrieux, K. Matković · Interactive Trajectories and Events Analysis: Exploring Sequences of Spatio-Temporal Events in Movement Data · *The Visual Computer, 32* · 2016

C.3 INVITED TALKS

- **L. Cibulski** · User Preferences in Visualization-Based Choice-Making · *IfI Colloquium* · Host: Prof. Jürgen Bernard · University of Zurich, Switzerland · 2023
- **L. Cibulski** · Making Choices With Visualizations · *CG Colloquy Cycle* · Host: Prof. Eduard Gröller · TU Wien, Austria · 2023
- **L. Cibulski** · Making Choices With Visualizations · *ICG Lab Talk Series* · Host: Prof. Marc Streit · Johannes Kepler University, Linz, Austria · 2023

C.4 REVIEWING ACTIVITIES

- *IEEE Transactions on Visualization and Computer Graphics* · 2024
- *IEEE Visualization Conference* · 2020 – 2024
- *Computer Graphics and Applications Special Issue on Visualization in the Wild* · 2023
- *Computer Graphics Forum* · 2022
- *IEEE Transactions on Evolutionary Computation* · 2021
- *Eurographics Conference on Visualization* · 2021
- *IEEE Conference on Visual Analytics Science and Technology* · 2019

D TEACHING EXPERIENCE

D.1 LECTURING

- L. Cibulski · Informationsdesign und Informationsvisualisierung · Course *User-Centered Design in Visual Computing* · Host: Prof. Jörn Kohlhammer · Technical University of Darmstadt, Germany · 2023

D.2 SUPERVISION OF MASTER THESES

- M. Langer · Nutzerorientierte Entwicklung einer Visualisierung für Messdaten aus der Fahrzeugerprobung · In cooperation with BMW · 2023
- T. Reiner · Visual Inference for Modeling and Reasoning With Bayesian Networks · 2023
- S. Hainzl · Bridging the Domain Gap – Visual Identification of Domain-Invariant Features in Time Series · 2022
- M. Nieslony · Visuelle Analyse des Antriebsdegradationsverhaltens von PHEV- und BEV-Fahrzeugen im Kontext der Dauerlauferprobung · In cooperation with BMW · 2021
- H. Pfeifer · A Visual Analytics Approach to Sensor Analysis for End-of-Line Testing · 2019

D.3 SUPERVISION OF INTERNSHIPS AND SEMINARS

- N. Lange · Visual Analysis for Multi-Attribute Choice · Seminar Visual Analytics · Summer 2023
- M. Bange · Visual Analysis for Multi-Attribute Choice · Seminar Visual Analytics · Summer 2022
- M. Bange, Y. Wu, Q. Yu · Import and Layout of XML Failure Models · Internship Visual Computing · Winter 2021
- A. Kruk · Visualizing States of Technical Systems · Seminar Visual Analytics · Summer 2021
- M. Stahr · Failure Classification of Electric Machines · Student Assistant · Jan - Sep 2020

- T. Dollenbacher · Integrated Display of Parametric 3D Models in Web-Based Parallel Coordinates · Internship Visual Computing · Summer 2020
- L. Schäfer · A Visual Programming Interface for Data Preparation Scripts · Internship Visual Computing · Summer 2020
- M. Pedko · Visual Analysis of Industrial Data · Seminar Visual Analytics · Summer 2020
- R. Garbe, A. Mulder, J. Thao · Processing User Interaction for Preference Modeling · Internship Visual Computing · Winter 2019
- C. Werner · Clustering Multivariate Data Using Self-Organizing Maps · Internship Visual Computing · Winter 2019
- D. Jeckel · Interactive Multiple-Criteria Decision-Making · Seminar Visual Analytics · Summer 2019
- C. Gonzalez · Visual Analysis Approaches to Anomaly Detection in Temporal Data · Seminar Visual Analytics · Summer 2018
- F. Otto · Visual Analysis Approaches to Time Series Prediction · Seminar Visual Analytics · Summer 2018

D.4 GRADUATE TEACHING ASSISTANT

- Introduction to Computer Science · Tutorial: Supporting first-semester students by recapitulating and deepening the lecture contents. Providing assistance in solving exercises. · Otto-von-Guericke University, Magdeburg · Oct 2013 – Oct 2014

E

CURRICULUM VITAE

LENA CIBULSKI

EDUCATION

Dez 2017 – Mar 2024 TECHNICAL UNIVERSITY, Darmstadt

Defense: 25th Mar 2024 · Focus on Visualization for Decision-Making — *Ph.D. in Computer Science*
Doctoral Thesis: *Visual Analysis for Multi-Attribute Choice*

Apr 2016 – Nov 2017 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Focus on Visual Analytics · Minor in Gearing Technology — *M.Sc. in Computer Science*
Master Thesis: *Towards Visual Feature Selection for Multivariate, Time-Oriented Data*

Oct 2012 – Apr 2016 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Focus on Computer-Assisted Medical Applications — *B.Sc. in Visual Computing*
Bachelor Thesis: *Interactive Visual Analysis of Gear Stress as Surface Ensemble*

WORK EXPERIENCE

Nov 2017 – Jun 2024 FRAUNHOFER IGD, Darmstadt

Execution and lead of industrial and research projects that aim at informing multi-attribute decisions in engineering by using interactive visualization. — *Researcher*
Department of Information Visualization and Visual Analytics

Jan 2023 – Mar 2023 JOHANNES KEPLER UNIVERSITY, Linz

Visualization support for the design of electric drives including multi-topology and multi-component aspects. Integration of PAVED into the engineering software. — *Visiting Researcher*
Institute for Electrical Drives and Power Electronics

Apr 2017 – Oct 2017 FRAUNHOFER IGD, Darmstadt

Exploration of relationships in multivariate, time-dependent data for the purpose of feature selection for regression analysis. — *Research Assistant Master Thesis*
Department of Information Visualization and Visual Analytics

Mar 2016 – Aug 2016 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Visual analytics concepts to support epidemiologists in generating hypotheses based on exploration of population study data. — *Research Assistant*
Visualization Group at the Institute of Simulation and Graphics

Oct 2015 – Mar 2016 VRVIS RESEARCH CENTER, Vienna

Visual analysis of spatio-temporal events data. Visualization techniques for the analysis of gear stress as data surface ensembles. — *Research Assistant Bachelor Thesis*
Interactive Visualization Group

Dec 2012 – Sep 2015 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Volunteer Helper Student Council of the Faculty of Computer Science

Mar 2015 – Aug 2015 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Research Assistant Visual-interactive filtering to support the analysis of cohort study data. Visualization Group at the Institute of Simulation and Graphics

Sep 2013 – Sep 2014 OTTO-VON-GUERICKE UNIVERSITY, Magdeburg

Organization Preliminary courses offered by the Faculty of Computer Science

CERTIFICATIONS AND AWARDS

Apr 2024 · Certificate Higher Education Teaching – Basic Competencies · Technical University, Darmstadt

Nov 2023 · Best Paper Honorable Mention "Impact on Science" · Fraunhofer-Institute for Computer Graphics Research IGD, Darmstadt

Apr 2019 · Fraunhofer TALENTA start · Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V., Munich

Apr 2019 · Certified ScrumMaster (CSM) · Scrum Alliance · Valid for 2 years

Dec 2016 · Rudolf-Kruse-Award · Faculty of Computer Science, Otto-von-Guericke University, Magdeburg

Apr 2016 · Best Graduate of the Academic Year in the Visual Computing Bachelor's Degree · Faculty of Computer Science, Otto-von-Guericke University, Magdeburg

Oct 2013 · Deutschland-Stipendium · Federal Ministry of Education and Research, Germany

BIBLIOGRAPHY

[1] Paola Adinolfi. "A Journey Around Decision-Making: Searching for the "Big Picture" Across Disciplines." In: *European Management Journal* 39.1 (2021), pp. 9–21.

[2] Christopher Ahlberg and Ben Shneiderman. "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1994, pp. 313–317.

[3] Jae-Wook Ahn and Peter Brusilovsky. "Adaptive Visualization for Exploratory Information Retrieval." In: *Information Processing & Management* 49.5 (2013), pp. 1139–1164.

[4] Wolfgang Aigner, Silvia Miksch, Heidrun Schumann, and Christian Tominski. *Visualization of Time-Oriented Data*. Second. Springer, 2023.

[5] Sara Alspaugh, Nava Zokaei, Andrea Liu, Cindy Jin, and Marti A. Hearst. "Futzing and Moseying: Interviews with Professional Data Analysts on Exploration Practices." In: *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), pp. 22–31.

[6] Tomás Alves, Tiago Delgado, Joana Henriques-Calado, Daniel Gonçalves, and Sandra Gama. "Exploring the Role of Conscientiousness on Visualization-Supported Decision-Making." In: *Computers & Graphics* 111 (2023), pp. 47–62.

[7] Robert A. Amar and John T. Stasko. "Knowledge Precepts for Design and Evaluation of Information Visualizations." In: *IEEE Transactions on Visualization and Computer Graphics* 11.4 (2005), pp. 432–442.

[8] Robert Amar, James Eagan, and John Stasko. "Low-Level Components of Analytic Activity in Information Visualization." In: *Proceedings of the IEEE Symposium on Information Visualization*. 2005, pp. 111–117.

[9] Keith Andrews, Majda Osmić, and Gerhard Schagerl. "Aggregated Parallel Coordinates: Integrating Hierarchical Dimensions into Parallel Coordinates Visualisations." In: *Proceedings of the International Conference on Knowledge Technologies and Data-Driven Business*. 2015, pp. 1–4.

[10] Gennady Andrienko and Natalia Andrienko. "Constructing Parallel Coordinates Plot for Problem Solving." In: *Proceedings of the International Symposium on Smart Graphics*. 2001, pp. 9–14.

[11] Gennady Andrienko, Natalia Andrienko, and Piotr Jankowski. "Building Spatial Decision Support Tools for Individuals and Groups." In: *Journal of Decision Systems* 12.2 (2003), pp. 193–208.

[12] Natalia Andrienko and Gennady Andrienko. "Informed Spatial Decisions Through Coordinated Views." In: *Information Visualization* 2.4 (2003), pp. 270–285.

[13] Clemens Arbesser, Thomas Mühlbacher, Stefan Komornyik, and Harald Piringer. "Visual Analytics for Domain Experts: Challenges and Lessons Learned." In: *Proceedings of the International Symposium on Virtual Reality & Visual Computing*. 2017.

[14] Bon A. Aseniero, Tiffany Wun, David Ledo, Guenther Ruhe, Anthony Tang, and Sheelagh Carpendale. "STRATOS: Using Visualization to Support Decisions in Strategic Software Release Planning." In: *Proceedings of the Annual ACM Conference on Human Factors in Computing Systems*. 2015, pp. 1479–1488.

[15] Roger Atkinson. "Project Management: Cost, Time and Quality, Two Best Guesses and a Phenomenon, Its Time to Accept Other Success Criteria." In: *International Journal of Project Management* 17.6 (1999), pp. 337–342.

[16] Miguel Bagajewicz and Enmanuel Cabrera. "Pareto Optimal Solutions Visualization Techniques for Multiobjective Design and Upgrade of Instrumentation Networks." In: *Industrial & Engineering Chemistry Research* 42.21 (2003), pp. 5195–5203.

[17] Sanjana Bajracharya, Giuseppe Carenini, Brent Chamberlain, Kai D. Chen, Daniel Klein, David Poole, Hamed Taheri, and Gunilla Öberg. "Interactive Visualization for Group Decision Analysis." In: *International Journal of Information Technology & Decision Making* 17.6 (2018), pp. 1839–1864.

[18] Aaron Bangor, Philip Kortum, and James Miller. "Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale." In: *Journal of Usability Studies* 4.3 (2009), pp. 114–123.

[19] Chester Barnard. *The Functions of the Executive*. Harvard University Press, 1938.

[20] Aaron Barsky, Tamara Munzner, Jennifer Gardy, and Robert Kincaid. "Cerebral: Visualizing Multiple Experimental Conditions on a Graph with Biological Context." In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1253–1260.

[21] Clara-Maria Barth, Jenny Schmid, Ibrahim Al Hazwani, Madhav Sachdeva, Lena Cibulski, and Jürgen Bernard. "How Applicable are Attribute-Based Approaches for Human-Centered Ranking Creation?" In: *Computers & Graphics* 114 (2023), pp. 45–58.

[22] Rahul C. Basole, Ahsan Qamar, Hyunwoo Park, Christiaan J. J. Paredis, and Leon F. McGinnis. "Visual Analytics for Early-Phase Complex Engineered System Design Support." In: *IEEE Computer Graphics and Applications* 35.2 (2015), pp. 41–51.

[23] Dominikus Baur, Frederik Seiffert, Michael Sedlmair, and Sebastian Boring. "The Streams of our Lives: Visualizing Listening Histories in Context." In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1119–1128.

[24] Jeanette Bautista and Giuseppe Carenini. "An Integrated Task-Based Framework for the Design and Evaluation of Visualizations to Support Preferential Choice." In: *Proceedings of the Working Conference on Advanced Visual Interfaces*. 2006, pp. 217–224.

[25] Lee R. Beach and Raanan Lipshitz. "Why Classical Decision Theory is an Inappropriate Standard for Evaluating and Aiding Most Human Decision Making." In: *Decision Making in Aviation* 85 (2017), pp. 835–847.

[26] Richard A. Becker and William S. Cleveland. "Brushing Scatterplots." In: *Technometrics* 29.2 (1987), pp. 127–142.

[27] Michael Beham, Wolfgang Herzner, Eduard Gröller, and Johannes Kehrer. "Cupid: Cluster-Based Exploration of Geometry Generators with Parallel Coordinates and Radial Trees." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1693–1702.

[28] Wolfgang Berger, Harald Piringer, Peter Filzmoser, and Eduard Gröller. "Uncertainty-Aware Exploration of Continuous Parameter Spaces Using Multivariate Prediction." In: *Computer Graphics Forum* 30.3 (2011), pp. 911–920.

[29] Steven Bergner, Michael Sedlmair, Torsten Möller, Sareh N. Abdolyousefi, and Ahmed Saad. "ParaGlide: Interactive Parameter Space Partitioning for Computer Simulations." In: *IEEE Transactions on Visualization and Computer Graphics* 19.9 (2013), pp. 1499–1512.

[30] Jacques Bertin. *Semiology of Graphics: Diagrams, Networks, Maps*. Esri Press, 1983.

[31] James R. Bettman, Mary F. Luce, and John W. Payne. "Constructive Consumer Choice Processes." In: *Journal of Consumer Research* 25.3 (1998), pp. 187–217.

[32] James R. Bettman and Michel A. Zins. “Constructive Processes in Consumer Choice.” In: *Journal of Consumer Research* 4.2 (1977), pp. 75–85.

[33] Diwas Bhattarai, Jian Zhang, and Bijaya B. Karki. “Parallel Coordinates-Based Visual Analytics for Materials Property.” In: *Proceedings of the International Conference on Information Visualization Theory and Applications*. 2019, pp. 83–95.

[34] Georges-Pierre Bonneau, Hans-Christian Hege, Chris R. Johnson, Manuel M. Oliveira, Kristin Potter, Penny Rheingans, and Thomas Schultz. “Overview and State-of-the-Art of Uncertainty Visualization.” In: *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization* (2014), pp. 3–27.

[35] Maryam Booshehrian, Torsten Möller, Randall M. Peterman, and Tamara Munzner. “Vismon: Facilitating Analysis of Trade-offs, Uncertainty, and Sensitivity in Fisheries Management Decision Making.” In: *Computer Graphics Forum* 31.3pt3 (2012), pp. 1235–1244.

[36] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. “D^3 Data-Driven Documents.” In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2301–2309.

[37] Matthew Brehmer, Stephen Ingram, Jonathan Stray, and Tamara Munzner. “Overview: The Design, Adoption, and Analysis of a Visual Document Mining Tool for Investigative Journalists.” In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2271–2280.

[38] Matthew Brehmer and Tamara Munzner. “A Multi-Level Typology of Abstract Visualization Tasks.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2376–2385.

[39] John Brooke. “SUS: A Quick and Dirty Usability Scale.” In: *Usability Evaluation in Industry* 189.3 (1996), pp. 189–194.

[40] Nathan Brown and Caitlin Mueller. “Designing with Data: Moving Beyond the Design Space Catalog.” In: *Proceedings of the Annual Conference of the Association for Computer Aided Design in Architecture: Disciplines and Disruption*. 2017, pp. 154–163.

[41] Jerome R. Busemeyer, Sebastian Gluth, Jörg Rieskamp, and Brandon M. Turner. “Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-Based Decisions.” In: *Trends in Cognitive Sciences* 23.3 (2019), pp. 251–263.

[42] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision To Think*. Morgan Kaufman Publishers, Inc., 1999.

[43] Giuseppe Carenini and John Loyd. "ValueCharts: Analyzing Linear Models Expressing Preferences and Evaluations." In: *Proceedings of the Working Conference on Advanced Visual Interfaces*. 2004, pp. 150–157.

[44] Susan Carlson-Skalak, Michael D. White, and Yong Teng. "Using an Evolutionary Algorithm for Catalog Design." In: *Research in Engineering Design* 10.2 (1998), pp. 63–83.

[45] Sheelagh Carpendale. "Evaluating Information Visualizations." In: *Information Visualization*. Springer, 2008, pp. 19–45.

[46] Davide Ceneda, Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jörg Schulz, Marc Streit, and Christian Tominski. "Characterizing Guidance in Visual Analytics." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), pp. 111–120.

[47] Mehdi R. Chakhchoukh, Nadia Boukhelifa, and Anastasia Bezerianos. "Understanding How In-Visualization Provenance Can Support Trade-off Analysis." In: *IEEE Transactions on Visualization and Computer Graphics* 29.9 (2022), pp. 3758–3774.

[48] Chaomei Chen. "Top 10 Unsolved Information Visualization Problems." In: *IEEE Computer Graphics and Applications* 25.4 (2005), pp. 12–16.

[49] Shahar Chen, David Amid, Ofer M. Shir, Lior Limonad, David Boaz, Ateret Anaby-Tavor, and Tobias Schreck. "Self-Organizing Maps for Multi-Objective Pareto Frontiers." In: *Proceedings of the IEEE Pacific Visualization Symposium*. 2013, pp. 153–160.

[50] Shenghui Cheng and Klaus Mueller. "The Data Context Map: Fusing Data and Attributes into a Unified Display." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2015), pp. 121–130.

[51] Lena Cibulski, Evanthia Dimara, Setia Hermawati, and Jörn Kohlhammer. "Supporting Domain Characterization in Visualization Design Studies With the Critical Decision Method." In: *Proceedings of the IEEE Workshop on Visualization Guidelines in Research, Design, and Education*. 2022.

[52] Lena Cibulski, Thorsten May, Johanna Schmidt, and Jörn Kohlhammer. "COMPO*SED: Composite Parallel Coordinates for Co-Dependent Multi-Attribute Choices." In: *IEEE Transactions on Visualization and Computer Graphics* 29.10 (2022), pp. 4047–4061.

[53] Lena Cibulski and Hubert Mitterhofer. "Interaktive Visualisierung: Durchblick beim Produktdesign." In: *WT Werkstattstechnik online* (2021), pp. 248–250.

[54] Lena Cibulski, Hubert Mitterhofer, Thorsten May, and Jörn Kohlhammer. "PAVED: Pareto Front Visualization for Engineering Design." In: *Computer Graphics Forum* 39.3 (2020), pp. 405–416.

[55] Jarry Claessen and Jarke J. van Wijk. "Flexible Linked Axes for Multivariate Data Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2310–2316.

[56] Edgar F. Codd. "A Relational Model of Data for Large Shared Data Banks." In: *Communications of the ACM* 13.6 (1970), 377–387.

[57] Cristina Conati, Giuseppe Carenini, Enamul Hoque, Ben Steichen, and Dereck Toker. "Evaluating the Impact of User Characteristics and Different Layouts on an Interactive Visualization for Decision Making." In: *Computer Graphics Forum* 33.3 (2014), pp. 371–380.

[58] Kalyanmoy Deb. "Multi-Objective Optimisation Using Evolutionary Algorithms: an Introduction." In: *Multi-Objective Evolutionary Optimisation for Product Design and Manufacturing*. Springer, 2011, pp. 3–34.

[59] *Decision Making*. In: *The American Psychological Association Dictionary of Psychology*. URL: https://dictionary.apa.org/decision-making (visited on 11/06/2023).

[60] Ravi Dhar. "New Directions in Behavioral Decision Theory: Implications for Consumer Choice." In: *Advances in Consumer Research* 22.1 (1995), p. 203.

[61] Luis C. Dias and Alexis Tsoukiàs. "On the Constructive and Other Approaches in Decision Aiding." In: *Proceedings of the Meeting of the European Working Group on Multiple Criteria Decision Aiding*. 2004, pp. 13–28.

[62] George E. Dieter and Linda C. Schmidt. *Engineering Design*. McGraw Hill, 2009.

[63] Evanthia Dimara, Anastasia Bezerianos, and Pierre Dragicevic. "The Attraction Effect in Information Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), pp. 471–480.

[64] Evanthia Dimara, Anastasia Bezerianos, and Pierre Dragicevic. "Conceptual and Methodological Issues in Evaluating Multidimensional Visualizations for Decision Support." In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2017), pp. 749–759.

[65] Evanthia Dimara, Steven Franconeri, Catherine Plaisant, Anastasia Bezerianos, and Pierre Dragicevic. "A Task-Based Taxonomy of Cognitive Biases for Information Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 26.2 (2018), pp. 1413–1432.

[66] Evanthia Dimara and Charles Perin. "What is Interaction for Data Visualization?" In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2019), pp. 119–129.

[67] Evanthia Dimara and John Stasko. "A Critical Reflection on Visualization Research: Where Do Decision Making Tasks Hide?" In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2021), pp. 1128–1138.

[68] Evanthia Dimara, Harry Zhang, Melanie Tory, and Steven Franconeri. "The Unmet Data Visualization Needs of Decision Makers within Organizations." In: *IEEE Transactions on Visualization and Computer Graphics* 28.12 (2021), pp. 4101–4112.

[69] Helmut Doleisch and Helwig Hauser. "Smooth Brushing for Focus+Context Visualization of Simulation Data in 3D." In: *Proceedings of the International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*. 2002, pp. 147–154.

[70] Yao Duan and Dan M. Ionel. "A Review of Recent Developments in Electrical Machine Design Optimization Methods with a Permanent-Magnet Synchronous Motor Benchmark Study." In: *IEEE Transactions on Industry Applications* 49.3 (2013), pp. 1268–1275.

[71] Bianchi Dy, Nazim Ibrahim, Ate Poorthuis, and Sam Joyce. "Improving Visualization Design for Effective Multi-Objective Decision Making." In: *IEEE Transactions on Visualization and Computer Graphics* 28.10 (2021), pp. 3405–3416.

[72] Joscha Eirich, Jakob Bonart, Dominik Jäckle, Michael Sedlmair, Ute Schmid, Kai Fischbach, Tobias Schreck, and Jürgen Bernard. "IRVINE: A design study on analyzing correlation patterns of electrical engines." In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2021), pp. 11–21.

[73] Niklas Elmqvist, Pierre Dragicevic, and Jean-Daniel Fekete. "Rolling the Dice: Multidimensional Visual Exploration Using Scatterplot Matrix Navigation." In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1539–1148.

[74] Accreditation Board for Engineering and Technolog (ABET). *Criteria for Accrediting Engineering Programs, 2021 - 2022*. https://www.abet.org/accreditation/accreditation-

criteria/criteria-for-accrediting-engineering-programs-2021-2022/. [Accessed 21 June 2021].

[75] Paolo Federico, Markus Wagner, Alexander Rind, Albert Amor-Amorós, Silvia Miksch, and Wolfgang Aigner. "The Role of Explicit Knowledge: A Conceptual Model of Knowledge-Assisted Visual Analytics." In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. 2017, pp. 92–103.

[76] Nivan Ferreira, Marcos Lage, Harish Doraiswamy, Huy Vo, Luc Wilson, Heidi Werner, Muchan Park, and Cláudio Silva. "Urbane: A 3D Framework to Support Data Driven Decision Making in Urban Development." In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. 2015, pp. 97–104.

[77] Peter J. Fleming, Robin C. Purshouse, and Robert J. Lygoe. "Many-Objective Optimization: An Engineering Design Perspective." In: *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*. 2005, pp. 14–32.

[78] Laura A. Garrison, Juliane Müller, Stefanie Schreiber, Steffen Oeltze-Jafra, Helwig Hauser, and Stefan Bruckner. "DimLift: Interactive Hierarchical Data Exploration through Dimensional Bundling." In: *IEEE Transactions on Visualization and Computer Graphics* 27.6 (2021), pp. 2908–2922.

[79] Rohan K. Gavade. "Multi-Criteria Decision Making: An Overview of Different Selection Problems and Methods." In: *International Journal of Computer Science and Information Technologies* 5.4 (2014), pp. 5643–5646.

[80] Amit Goel, Chuck Baker, Clifford A. Shaffer, Bernard Grossman, Raphael T. Haftka, William H. Mason, and Layne T. Watson. "VizCraft: A Multidimensional Visualization Tool for Aircraft Configuration Design." In: *Proceedings of the IEEE Conference on Visualization*. 1999, pp. 425–428.

[81] Victor González and Alfred Kobsa. "A Workplace Study of the Adoption of Information Visualization Systems." In: *Proceedings of the International Conference on Knowledge Management*. 2003, pp. 92–102.

[82] Sarah Goodwin, Jason Dykes, Sara Jones, Iain Dillingham, Graham Dove, Alison Duffy, Alexander Kachkaev, Aidan Slingsby, and Jo Wood. "Creative User-Centered Visualization Design for Energy Analysts and Modelers." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2516–2525.

[83] Martin Graham and Jessie Kennedy. "Using Curves to Enhance Parallel Coordinate Visualisations." In: *Proceedings of the International Conference on Information Visualization*. 2003, pp. 10–16.

[84] Samuel Gratzl, Nils Gehlenborg, Alexander Lex, Hanspeter Pfister, and Marc Streit. "Domino: Extracting, Comparing, and Manipulating Subsets Across Multiple Tabular Datasets." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2023–2032.

[85] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. "Lineup: Visual Analysis of Multi-Attribute Rankings." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2277–2286.

[86] Rebecca A. Grier, Aaron Bangor, Philip Kortum, and S. Camille Peres. "The System Usability Scale: Beyond Standard Usability Testing." In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 57. 1. 2013, pp. 187–191.

[87] Shunan Guo, Fan Du, Sana Malik, Eunyee Koh, Sungchul Kim, Zhicheng Liu, Donghyun Kim, Hongyuan Zha, and Nan Cao. "Visualizing Uncertainty and Alternatives in Event Sequence Predictions." In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2019, pp. 1–12.

[88] Jussi Hakanen, Kaisa Miettinen, and Krešimir Matković. "Task-Based Visual Analytics for Interactive Multiobjective Optimization." In: *Journal of the Operational Research Society* 72.9 (2021), pp. 2073–2090.

[89] Jussi Hakanen, Sanjin Radoš, Giovanni Misitano, Bhupinder S. Saini, Kaisa Miettinen, and Krešimir Matković. "Interactivized: Visual Interaction for Better Decisions with Interactive Multiobjective Optimization." In: *IEEE Access* 10 (2022), pp. 33661–33678.

[90] Robert Harris. *Introduction to Decision Making*. Virtual Salt. `http://www.virtualsalt.com/crebook5.htm`. 1998.

[91] Harry P. Hatry and Kathryn E. Newcomer. "Pitfalls in Evaluations." In: *Handbook of Practical Program Evaluation* (2015), pp. 701–724.

[92] Helwig Hauser. "Generalizing Focus+Context Visualization." In: *Scientific Visualization: The Visual Extraction of Knowledge from Data*. Springer, 2006, pp. 305–327.

[93] George A. Hazelrigg. "A Framework for Decision-Based Engineering Design." In: *Journal of Mechanical Design* 120 (1998), pp. 653–658.

[94] Jeffrey Heer and Ben Shneiderman. "Interactive Dynamics for Visual Analysis." In: *Queue* 10.2 (2012), pp. 30–55.

[95] Julian Heinrich and Daniel Weiskopf. "Continuous Parallel Coordinates." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1531–1538.

[96] Julian Heinrich and Daniel Weiskopf. "State of the Art of Parallel Coordinates." In: *Eurographics - State of the Art Reports*. 2013, pp. 95–116.

[97] Setia Hermawati, Lena Cibulski, and Glyn Lawson. "Towards Using the Critical Decision Method for Studying Visualisation-Based Decision-Making." In: *Contemporary Ergonomics & Human Factors*. Chartered Institute of Ergonomics & Human Factors, 2021, pp. 130–136.

[98] Jan Hettenhausen, Andrew Lewis, and Sanaz Mostaghim. "Interactive Multi-Objective Particle Swarm Optimization With Heatmap-Visualization-Based User Interface." In: *Engineering Optimization* 42.2 (2010), pp. 119–139.

[99] Emily Hindalong, Jordon Johnson, Giuseppe Carenini, and Tamara Munzner. "Towards Rigorously Designed Preference Visualizations for Group Decision Making." In: *Proceedings of the IEEE Pacific Visualization Symposium*. 2020, pp. 181–190.

[100] Patrick Hoffman, Georges Grinstein, Kenneth Marx, Ivo Grosse, and Eugene Stanley. "DNA Visual and Analytic Data Mining." In: *Proceedings of the IEEE Conference on Visualization*. 1997, pp. 437–441.

[101] Robert R. Hoffman, Beth Crandall, and Nigel Shadbolt. "Use of the Critical Decision Method to Elicit Expert Knowledge: A Case Study in the Methodology of Cognitive Task Analysis." In: *Human Factors* 40.2 (1998), pp. 254–276.

[102] Robin M. Hogarth. *Judgement and Choice: The Psychology of Decision*. John Wiley & Sons, 1987.

[103] Jeffrey Horn. "Multicriteria Decision Making and Evolutionary Computation." In: *Handbook of Evolutionary Computation*. Institute of Physics Publishing, 1997.

[104] Dandan Huang, Melanie Tory, Bon A. Aseniero, Lyn Bartram, Scott Bateman, Sheelagh Carpendale, Anthony Tang, and Robert Woodbury. "Personal Visualization and Personal Visual Analytics." In: *IEEE Transactions on Visualization and Computer Graphics* 21.3 (2014), pp. 420–433.

[105] Yansong Huang, Zherui Zhang, Ao Jiao, Yuxin Ma, and Ran Cheng. "A Comparative Visual Analytics Framework for Evaluating Evolutionary Processes in Multi-objective Optimization." In: *IEEE Transactions on Visualization and Computer Graphics* 30.1 (2023), pp. 661–671.

[106] Christian Hurson and Constantin Zopounidis. "On the Use of Multicriteria Decision Aid Methods to Portfolio Selection." In: *Proceedings of the International Conference on Multicriteria Analysis*. 1997, pp. 496–507.

[107] ISO. *ISO 9241-210:2019 Human-centered design for interactive systems*. 2019.

[108] International Technology Education Association (ITEA). *Standards for Technological Literacy: Content for the Study of Technology*. 2000.

[109] Alfred Inselberg. "The Plane with Parallel Coordinates." In: *The Visual Computer* 1.2 (1985), pp. 69–91.

[110] Alfred Inselberg and Bernard Dimsdale. "Parallel Coordinates: A Tool for Visualizing Multi-Dimensional Geometry." In: *Proceedings of the IEEE Conference on Visualization*. 1990, pp. 361–378.

[111] Petra Isenberg, Anthony Tang, and Sheelagh Carpendale. "An Exploratory Study of Visual Information Analysis." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2008, pp. 1217–1226.

[112] Petra Isenberg, Torre Zuk, Christopher Collins, and Sheelagh Carpendale. "Grounded Evaluation of Information Visualizations." In: *Proceedings of the Workshop on BEyond time and errors: novel evaLuation methods for Information Visualization*. 2008, pp. 1–8.

[113] Tobias Isenberg, Petra Isenberg, Jian Chen, Michael Sedlmair, and Torsten Möller. "A Systematic Review on the Practice of Evaluating Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2818–2827.

[114] Alessio Ishizaka and Philippe Nemery. *Multi-Criteria Decision Analysis: Methods and Software*. John Wiley & Sons, 2013.

[115] Waqas Javed and Niklas Elmqvist. "Exploring the Design Space of Composite Visualization." In: *Proceedings of the IEEE Pacific Visualization Symposium*. 2012, pp. 1–8.

[116] Yu Jin, Laura Koesten, and Torsten Möller. "Exploring the Design Space of Three Criteria Decision Making." In: *IEEE Computer Graphics and Applications* 43.5 (2023), pp. 26–38.

[117] Jimmy Johansson and Camilla Forsell. "Evaluation of Parallel Coordinates: Overview, Categorization and Guidelines for Future Research." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2015), pp. 579–588.

[118] Richard M. Johnson. "Trade-Off Analysis of Consumer Values." In: *Journal of Marketing Research* 11.2 (1974), pp. 121–127.

[119] Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

[120] Makbule Kandakoglu, Grit Walther, and Sarah Ben Amor. "The Use of Multi-Criteria Decision-Making Methods in Project Portfolio Selection: A Literature Review and Future Research Directions." In: *Annals of Operations Research* (2023), pp. 1–24.

[121] Youn-ah Kang and John Stasko. "Examining the Use of a Visual Analytics System for Sensemaking Tasks: Case Studies with Domain Experts." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2869–2878.

[122] Johannes Kehrer and Helwig Hauser. "Visualization and Visual Analysis of Multifaceted Scientific Data: A Survey." In: *IEEE Transactions on Visualization and Computer Graphics* 19.3 (2012), pp. 495–513.

[123] Johannes Kehrer, Philipp Muigg, Helmut Doleisch, and Helwig Hauser. "Interactive Visual Analysis of Heterogeneous Scientific Data Across an Interface." In: *IEEE Transactions on Visualization and Computer Graphics* 17.7 (2010), pp. 934–946.

[124] Andreas Kerren and Falk Schreiber. "Toward the Role of Interaction in Visual Analytics." In: *Proceedings of the Winter Simulation Conference*. 2012, pp. 1–13.

[125] Robert Kincaid, Amir Ben-Dor, and Zohar Yakhini. "Exploratory Visualization of Array-Based Comparative Genomic Hybridization." In: *Information Visualization* 4.3 (2005), pp. 176–190.

[126] Timoleon Kipouros, Alfred Inselberg, Geoffrey Parks, and Mark Savill. "Parallel Coordinates in Computational Engineering Design." In: *Proceedings of the AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. 2013, p. 1750.

[127] David Kirsh. "Interactivity and Multimedia Interfaces." In: *Instructional Science* 25 (1997), pp. 79–96.

[128] Gary A. Klein. *Sources of Power: How People Make Decisions*. MIT Press, 2017.

[129] Gary A. Klein, Roberta Calderwood, and Donald MacGregor. "Critical Decision Method for Eliciting Knowledge." In: *IEEE Transactions on Systems, Man, and Cybernetics* 19.3 (1989), pp. 462–472.

[130] Gary A. Klein, Judith Orasanu, Roberta Calderwood, and Caroline E. Zsambok, eds. *Decision Making in Action: Models and Methods*. Ablex Publishing Coporation, 1993.

[131] Danny S. M. Koh, Hock G. Tee, Boon K. Soh, and Angela L. S. Tan. “Tools for Facilitating Critical Decision Method during Tacit Knowledge Elicitation.” In: *Proceedings of the International Conference on Naturalistic Decision Making*. 2015.

[132] Teuvo Kohonen. *Self-Organizing Maps*. Springer, 2001.

[133] Zoltán Konyha, Krešimir Matković, and Helwig Hauser. “Interactive Visual Analysis in Engineering: A Survey.” In: *Proceedings of the Spring Conference on Computer Graphics - Posters*. 2009, pp. 31–38.

[134] Pekka Korhonen, Herbert Moskowitz, and Jyrki Wallenius. “Multiple Criteria Decision Support: A Review.” In: *European Journal of Operational Research* 63.3 (1992), pp. 361–375.

[135] Pekka Korhonen and Jyrki Wallenius. “Visualization in the Multiple Objective Decision-Making Framework.” In: *Multiobjective Optimization*. Springer, 2008, pp. 195–212.

[136] Robert Kosara, Fabian Bendix, and Helwig Hauser. “Parallel Sets: Interactive Exploration and Visual Analysis of Categorical Data.” In: *IEEE Transactions on Visualization and Computer Graphics* 12.4 (2006), pp. 558–568.

[137] Alexander Kossiakoff, William N. Sweet, Samuel J. Seymour, and Steven M. Biemer. *Systems Engineering: Principles and Practices*. Second. John Wiley & Sons, 2011.

[138] Heidi Lam, Enrico Bertini, Petra Isenberg, Catherine Plaisant, and Sheelagh Carpendale. “Empirical Studies in Information Visualization: Seven Scenarios.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.9 (2011), pp. 1520–1536.

[139] Heidi Lam, Daniel Russell, Diane Tang, and Tamara Munzner. “Session Viewer: Visual Exploratory Analysis of Web Session Logs.” In: *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*. 2007, pp. 147–154.

[140] Ann Langley. “Between ’Paralysis by Analysis’ and ’Extinction by Instinct’.” In: *MIT Sloan Management Review* 36.3 (1995), p. 63.

[141] Tom Lanning, Kent Wittenburg, Michael Heinrichs, Christina Fyock, and Glenn Li. “Multidimensional Information Visualization Through Sliding Rods.” In: *Proceedings of the Working Conference on Advanced Visual Interfaces*. 2000, pp. 173–180.

[142] Yann Laurillau, Van-Bao Nguyen, Joëlle Coutaz, Gaëlle Calvary, Nadine Mandran, Fatoumata Camara, and Raffaella Balzarini. “The TOP-Slider for Multi-Criteria Decision Making by Non-Specialists.” In: *Proceedings of the Nordic Conference on Human-Computer Interaction*. 2018, pp. 642–653.

[143] Sukwon Lee, Sung-Hee Kim, and Bum C. Kwon. "VLAT: Development of a Visualization Literacy Assessment Test." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), pp. 551–560.

[144] James R. Lewis and Jeff Sauro. "The Factor Structure of the System Usability Scale." In: *Proceedings of the International Conference on Human Centered Design*. 2009, pp. 94–103.

[145] Guozheng Li, Runfei Li, Zicheng Wang, Chi H. Liu, Min Lu, and Guoren Wang. "HiTailor: Interactive Transformation and Visualization for Hierarchical Tabular Data." In: *IEEE Transactions on Visualization and Computer Graphics* 29.1 (2022), pp. 139–148.

[146] Jing Li, Jean-Bernard Martens, and Jarke J. van Wijk. "Judging Correlation from Scatterplots and Parallel Coordinate Plots." In: *Information Visualization* 9.1 (2010), pp. 13–30.

[147] Mats Lind, Jimmy Johansson, and Matthew Cooper. "Many-to-Many Relational Parallel Coordinates Displays." In: *Proceedings of the International Conference on Information Visualisation*. 2009, pp. 25–31.

[148] Claudio D. G. Linhares, Daniel M. Lima, Jean R. Ponciano, Mauro M. Olivatto, Marco A. Gutierrez, Jorge Poco, Caetano Traina, and Agma J. M. Traina. "ClinicalPath: A Visualization Tool to Improve the Evaluation of Electronic Health Records in Clinical Decision-Making." In: *IEEE Transactions on Visualization and Computer Graphics* 29.10 (2022), pp. 4031–4046.

[149] Ran Liu, Tiffany Chao, Catherine Plaisant, and Ben Shneiderman. *ManyLists: Product Comparison Tool Using Spatial Layouts with Animated Transitions*. Tech. rep. University of Maryland, 2012.

[150] Zhicheng Liu, Shamkant B. Navathe, and John T. Stasko. "Network-Based Visual Analysis of Tabular Data." In: *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*. 2011, pp. 41–50.

[151] David Lloyd and Jason Dykes. "Human-Centered Approaches in Geovisualization Design: Investigating Multiple Methods Through a Long-Term Case Study." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2498–2507.

[152] Alexander V. Lotov and Kaisa Miettinen. "Visualizing the Pareto Frontier." In: *Multiobjective Optimization*. Springer, 2008, pp. 213–243.

[153] Mary F. Luce, John W. Payne, and James R. Bettman. "Emotional Trade-Off Difficulty and Choice." In: *Journal of Marketing Research* 36.2 (1999), pp. 143–159.

[154] R. Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. Dover Publications, Inc., 2012.

[155] Yan Lyu, Fan Gao, I-Shuen Wu, and Brian Y. Lim. "Imma Sort by Two or More Attributes With Interpretable Monotonic Multi-Attribute Sorting." In: *IEEE Transactions on Visualization and Computer Graphics* 27.4 (2020), pp. 2369–2384.

[156] Katherine Malan and Gary Marsden. "Extending Dynamic Queries to Handle Uncertain Data." In: *South African Computer Journal* 37 (2006), pp. 61–74.

[157] Maayan S. Malter, Morris B. Holbrook, Barbara E. Kahn, Jeffrey R. Parker, and Donald R. Lehmann. "The Past, Present, and Future of Consumer Research." In: *Marketing Letters* 31 (2020), pp. 137–149.

[158] G. Elisabeta Marai. "Activity-Centered Domain Characterization for Problem-Driven Scientific Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2017), pp. 913–922.

[159] G. Elisabeta Marai and Torsten Möller. "The Fabric of Visualization." In: *Foundations of Data Visualization*. Springer, 2020, pp. 5–14.

[160] Alex Marsaw, Brian German, Peter Hollingsworth, Dimitri Mavris, and Steve Csonka. "An Interactive Visualization Environment for Decision Making in Aircraft Engine Preliminary Design." In: *Proceedings of the AIAA Aerospace Sciences Meeting and Exhibit*. 2007, p. 1333.

[161] Edmund Marth, Gerald Jungmayr, Wolfgang Amrhein, and Frank Jeske. "Magnetic-Geared Motor in Side-by-Side Arrangement-Optimization." In: *Proceedings of the IEEE International Electric Machines and Drives Conference*. 2019, pp. 839–846.

[162] Justin Matejka, Michael Glueck, Erin Bradner, Ali Hashemi, Tovi Grossman, and George Fitzmaurice. "Dream Lens: Exploration and Visualization of Large-Scale Generative Design Datasets." In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–12.

[163] Krešimir Matković, Mario Jelovic, Josip Juric, Zoltan Konyha, and Denis Gracanin. "Interactive Visual Analysis and Exploration of Injection Systems Simulations." In: *Proceedings of the IEEE Conference on Visualization*. 2005, pp. 391–398.

[164] Sean McKenna, Dominika Mazur, James Agutter, and Miriah Meyer. "Design Activity Framework for Visualization Design." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2191–2200.

[165] Matt McKeon. "Harnessing the Information Ecosystem with Wiki-Based Visualization Dashboards." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1081–1088.

[166] Peter McLachlan, Tamara Munzner, Eleftherios Koutsofios, and Stephen North. "LiveRAC: Interactive Visual Exploration of System Management Time-Series Data." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2008, pp. 1483–1492.

[167] Miriah Meyer and Jason Dykes. "Reflection on Reflection in Applied Visualization Research." In: *IEEE Computer Graphics and Applications* 38.6 (2018), pp. 9–16.

[168] Miriah Meyer and Jason Dykes. "Criteria for Rigor in Visualization Design Study." In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2019), pp. 87–97.

[169] Miriah Meyer, Tamara Munzner, and Hanspeter Pfister. "MizBee: A Multiscale Synteny Browser." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 897–904.

[170] Miriah Meyer, Tamara Munzner, Angela dePace, and Hanspeter Pfister. "MulteeSum: A Tool for Exploring Space-Time Expression Data." In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 908–917.

[171] Miriah Meyer, Michael Sedlmair, P. Samuel Quinan, and Tamara Munzner. "The Nested Blocks and Guidelines Model." In: *Information Visualization* 14.3 (2015), pp. 234–249.

[172] Miriah Meyer, Bang Wong, Mark Styczynski, Tamara Munzner, and Hanspeter Pfister. "Pathline: A Tool for Comparative Functional Genomics." In: *Computer Graphics Forum* 29.3 (2010), pp. 1043–1052.

[173] Kaisa Miettinen. *Nonlinear Multiobjective Optimization*. Springer Science & Business Media, 1999.

[174] Kaisa Miettinen. "Introduction to Multiobjective Optimization: Noninteractive Approaches." In: *Multiobjective Optimization*. Springer, 2008, pp. 1–26.

[175] Kaisa Miettinen. "Survey of Methods to Visualize Alternatives in Multiple Criteria Decision Making Problems." In: *OR Spectrum* 36.1 (2014), pp. 3–37.

[176] Silvia Miksch and Wolfgang Aigner. "A Matter of Time: Applying a Data–Users–Tasks Design Triangle to Visual Analytics of Time-Oriented Data." In: *Computers & Graphics* 38 (2014), pp. 286–290.

[177] Silvia Miksch, Heike Leitte, and Min Chen. "Knowledge-Assisted Visualization and Guidance." In: *Foundations of Data Visualization*. Springer, 2020, pp. 61–85.

[178] Laura G. Militello and Robert J. B. Hutton. "Applied Cognitive Task Analysis (ACTA): a Practitioner's Toolkit for Understanding Cognitive Task Demands." In: *Ergonomics* 41.11 (1998), pp. 1618–1641.

[179] Sanjay Mittal and Felix Frayman. "Towards a Generic Model of Configuration Tasks." In: *Proceedings of the International Joint Conference on Artificial Intelligence*. 1989, pp. 1395–1401.

[180] Caitlin T. Mueller and John A. Ochsendorf. "Combining Structural Performance and Designer Preferences in Evolutionary Design Space Exploration." In: *Automation in Construction* 52 (2015), pp. 70–82.

[181] Thomas Mühlbacher, Lorenz Linhardt, Torsten Möller, and Harald Piringer. "TreePOD: Sensitivity-Aware Selection of Pareto-Optimal Decision Trees." In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2017), pp. 174–183.

[182] Thomas Mühlbacher and Harald Piringer. "A Partition-Based Framework for Building and Validating Regression Models." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 1962–1971.

[183] Heinrich Müller et al. "Intuitive Visualization and Interactive Analysis of Pareto Sets Applied on Production Engineering System." In: *Success in Evolutionary Computation*. Springer, 2008, pp. 189–214.

[184] Juliane Müller, Mario Cypko, Alexander Oeser, Matthäus Stoehr, Veit Zebralla, Stefanie Schreiber, Susanne Wiegand, Andreas Dietz, and Steffen Oeltze-Jafra. "Visual Assistance in Clinical Decision Support." In: *Proceedings of the Eurographics Conference on Visualization - Dirk Bartz Prize*. 2021.

[185] Tamara Munzner. "A Nested Model for Visualization Design and Validation." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 921–928.

[186] Tamara Munzner. *Visualization Analysis and Design*. CRC Press, 2014.

[187] Tamara Munzner and Stephen North. *InfoVis 2003 Call for Participation: Papers*. https://www.infovis.org/infovis/2003/CFP/#papers. [Accessed 3 August 2023]. 2003.

[188] Dennis C. Neale, John M. Carroll, and Mary B. Rosson. "Evaluating Computer-Supported Cooperative Work: Models and Frameworks." In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. 2004, pp. 112–121.

[189] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2007.

[190] Chris North and Ben Shneiderman. "Snap-Together Visualization: A User Interface for Coordinating Visualizations via Relational Schemata." In: *Working Conference on Advanced Visual Interfaces*. 2000, pp. 128–135.

[191] Shigeru Obayashi, Shinkyu Jeong, and Kazuhisa Chiba. "Multi-Objective Design Exploration for Aerodynamic Configurations." In: *Proceedings of the AIAA Fluid Dynamics Conference and Exhibit*. 2005, p. 4666.

[192] Shigeru Obayashi and Daisuke Sasaki. "Visualization and Data Mining of Pareto Solutions Using Self-Organizing Map." In: *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*. 2003, pp. 796–809.

[193] Maria C. F. de Oliveira and Haim Levkowitz. "From Visual Data Exploration to Visual Data Mining: A Survey." In: *IEEE Transactions on Visualization and Computer Graphics* 9.3 (2003), pp. 378–394.

[194] Michael Oppermann and Tamara Munzner. "Data-First Visualization Design Studies." In: *Proceedings of the IEEE Workshop on Evaluation and Beyond - Methodological Approaches to Visualization*. 2020, pp. 74–80.

[195] Emre Oral, Ria Chawla, Michel Wijkstra, Narges Mahyar, and Evanthia Dimara. "From Information to Choice: A Critical Inquiry Into Visualization Tools for Decision Making." In: *IEEE Transactions on Visualization and Computer Graphics* 30.1 (2023), pp. 359–369.

[196] Judith Orasanu and Terry Connolly. "The Reinvention of Decision Making." In: *Decision Making in Action: Models and Methods*. Ablex Publishing Coporation, 1993, pp. 3–20.

[197] Lace M. Padilla, Sarah H. Creem-Regehr, Mary Hegarty, and Jeanine K. Stefanucci. "Decision Making with Visualizations: A Cognitive Framework Across Disciplines." In: *Cognitive Research: Principles and Implications* 3.1 (2018), pp. 1–25.

[198] Stephan Pajer, Marc Streit, Thomas Torsney-Weir, Florian Spechtenhauser, Torsten Möller, and Harald Piringer. "Weightlifter: Visual Weight Space Exploration for Multi-Criteria Decision Making." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), pp. 611–620.

[199] Paul Parsons. "Understanding Data Visualization Design Practice." In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2021), pp. 665–675.

[200] Paul Parsons, Colin M. Gray, Ali Baigelenov, and Ian Carr. "Design Judgment in Data Visualization Practice." In: *Proceedings of the IEEE Visualization Conference*. 2020, pp. 176–180.

[201] John W. Payne. "Task Complexity and Contingent Processing in Decision Making: An information Search and Protocol Analysis." In: *Organizational Behavior and Human Performance* 16.2 (1976), pp. 366–387.

[202] John W. Payne and James R. Bettman. "Preferential Choice and Adaptive Strategy Use." In: *Bounded Rationality: The Adaptive Toolbox*. MIT Press, 2001, pp. 123–146.

[203] John W. Payne, James R. Bettman, and Eric J. Johnson. *The Adaptive Decision Maker*. Cambridge University Press, 1993.

[204] Roy van Pelt, Rocco Gasteiger, Kai Lawonn, Monique Meuschke, and Bernhard Preim. "Comparative Blood Flow Visualization for Cerebral Aneurysm Treatment Assessment." In: *Computer Graphics Forum* 33.3 (2014), pp. 131–140.

[205] Adam Perer and Ben Shneiderman. "Integrating Statistics and Visualization: Case Studies of Gaining Clarity During Exploratory Data Analysis." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2008, pp. 265–274.

[206] Charles Perin, Pierre Dragicevic, and Jean-Daniel Fekete. "Revisiting Bertin Matrices: New Interactions for Crafting Tabular Visualizations." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2082–2091.

[207] Martin Pinzger, Harald Gall, Michael Fischer, and Michele Lanza. "Visualizing Multiple Evolution Metrics." In: *Proceedings of the ACM Symposium on Software Visualization*. 2005, pp. 67–75.

[208] Wiktor Piotrowski, Timoleon Kipouros, and P. John Clarkson. "Enhanced Interactive Parallel Coordinates Using Machine Learning and Uncertainty Propagation for Engineering Design." In: *Proceedings of the International Conference on eScience*. 2019, pp. 339–348.

[209] Harald Piringer, Wolfgang Berger, and Jürgen Krasser. "HyperMoVal: Interactive Visual Validation of Regression Models for Real-Time Simulation." In: *Computer Graphics Forum* 29.3 (2010), pp. 983–992.

[210] Peter Pirolli and Stuart Card. "The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis." In: *Proceedings of the International Conference on Intelligence Analysis*. 2005, pp. 2–4.

[211] Timothy J. Pleskac. "Decision and Choice: Luce's Choice Axiom." In: *International Encyclopedia of the Social & Behavioral Sciences* 5 (2015), pp. 895–900.

[212] Michael Polanyi. "The Tacit Dimension." In: *Knowledge in Organizations*. Routledge, 2009, pp. 135–146.

[213] Kristin Potter, Andrew Wilson, Peer-Timo Bremer, Dean Williams, Charles Doutriaux, Valerio Pascucci, and Chris R. Johnson. "Ensemble-Vis: A Framework for the Statistical Visualization of Ensemble Data." In: *Proceedings of the IEEE International Conference on Data Mining Workshops*. 2009, pp. 233–240.

[214] Zachary Pousman, John Stasko, and Michael Mateas. "Casual Information Visualization: Depictions of Data in Everyday Life." In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1145–1152.

[215] A. Johannes Pretorius, Mark-Anthony P. Bray, Anne E. Carpenter, and Roy A. Ruddle. "Visualization of Parameter Space for Image Analysis." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2402–2411.

[216] Andy Pryke, Sanaz Mostaghim, and Alireza Nazemi. "Heatmap Visualization of Population Based Multi Objective Algorithms." In: *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*. 2007, pp. 361–375.

[217] Pearl H. Z. Pu and Pratyush Kumar. "Evaluating Example-Based Search Tools." In: *Proceedings of the ACM Conference on Electronic Commerce*. 2004, pp. 208–217.

[218] Pearl Pu and Boi Faltings. "Enriching Buyers' Experiences: the SmartClient Approach." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2000, pp. 289–296.

[219] Robin C. Purshouse and Peter J. Fleming. "Conflict, Harmony, and Independence: Relationships in Evolutionary Multi-Criterion Optimisation." In: *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*. 2003, pp. 16–30.

[220] Antonio Rangel, Colin Camerer, and P. Read Montague. "A Framework for Studying the Neurobiology of Value-Based Decision Making." In: *Nature Reviews Neuroscience* 9.7 (2008), pp. 545–556.

[221] Patrick Riehmann, Jens Opolka, and Bernd Froehlich. "The Product Explorer: Decision Making with Ease." In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*. 2012, pp. 423–432.

[222] Patrick Riehmann, Andreas Schädler, Jannis Harder, Jakob Herpel, and Bernd Froehlich. “Configuration Finder: A Tidy Visual Interface for Effective Faceted Search.” In: *Proceedings of the Eurographics Conference on Visualization - Short Papers*. 2020, pp. 19–23.

[223] Horst W. J. Rittel and Melvin M. Webber. “Dilemmas in a General Theory of Planning.” In: *Policy Sciences* 4.3 (1973), pp. 155–169.

[224] Jonathan C. Roberts. “On Encouraging Multiple Views for Visualization.” In: *Proceedings of the International Conference on Information Visualization*. 1998, pp. 8–14.

[225] Jonathan C. Roberts. “State of the Art: Coordinated & Multiple Views in Exploratory Visualization.” In: *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*. 2007, pp. 61–71.

[226] Jonathan Roberts, Daniel A. Keim, Timothy Hanratty, Robert Rwolingson, Rick Walker, Mark Hall, Zack Jackobson, Valerie Lavigne, Chris Rooney, and Margaret Varga. “From Ill-Defined Problems to Informed Decisions.” In: *Proceedings of the International EuroVis Workshop on Visual Analytics*. 2014, pp. 7–11.

[227] Steven F. Roth and Joe Mattis. “Data Characterization for Intelligent Graphics Presentation.” In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1990, pp. 193–200.

[228] J. Edward Russo and Barbara A. Dosher. “Strategies for Multiattribute Binary Choice.” In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9.4 (1983), p. 676.

[229] Dominik Sacha, Andreas Stoffel, Florian Stoffel, Bum C. Kwon, Geoffrey Ellis, and Daniel A Keim. “Knowledge Generation Model for Visual Analytics.” In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1604–1613.

[230] Andrea Saltelli. “Sensitivity Analysis for Importance Assessment.” In: *Risk Analysis* 22.3 (2002), pp. 579–590.

[231] Purvi Saraiya, Chris North, Vy Lam, and Karen A. Duca. “An Insight-Based Longitudinal Study of Visual Analytics.” In: *IEEE Transactions on Visualization and Computer Graphics* 12.6 (2006), pp. 1511–1522.

[232] Jeff Sauro. *Measuring Usability with the System Usability Scale (SUS)*. https://measuringu.com/sus. [Accessed 23 July 2023]. 2011.

[233] Jenny Schmid and Jürgen Bernard. “A Taxonomy of Attribute Scoring Functions.” In: *Proceedings of the International EuroVis Workshop on Visual Analytics*. 2021, pp. 31–35.

[234] Jenny Schmid, Lena Cibulski, Ibrahim Al Hazwani, and Jürgen Bernard. "RankASco: A Visual Analytics Approach to Leverage Attribute-Based User Preferences for Item Rankings." In: *Proceedings of the International EuroVis Workshop on Visual Analytics*. 2022, pp. 7–11.

[235] Hans-Jörg Schulz, Thomas Nocke, Magnus Heitzler, and Heidrun Schumann. "A Design Space of Visualization Tasks." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2366–2375.

[236] Michael Sedlmair. "Design Study Contributions Come in Different Guises: Seven Guiding Scenarios." In: *Proceedings of the Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*. 2016, pp. 152–161.

[237] Michael Sedlmair, Christian Bernhold, Daniel Herrscher, Sebastian Boring, and Andreas Butz. "MostVis: An Interactive Visualization Supporting Automotive Engineers in Most Catalog Exploration." In: *Proceedings of the International Conference on Information Visualisation*. 2009, pp. 173–182.

[238] Michael Sedlmair, Matt Brehmer, Stephen Ingram, and Tamara Munzner. *Dimensionality Reduction in the Wild: Gaps and Guidance*. Tech. rep. University of British Columbia, 2012.

[239] Michael Sedlmair, Annika Frank, Tamara Munzner, and Andreas Butz. "RelEx: Visualization for Actively Changing Overlay Network Specifications." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2729–2738.

[240] Michael Sedlmair, Christoph Heinzl, Stefan Bruckner, Harald Piringer, and Torsten Möller. "Visual Parameter Space Analysis: A Conceptual Framework." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2161–2170.

[241] Michael Sedlmair, Petra Isenberg, Dominikus Baur, Michael Mauerer, Christian Pigorsch, and Andreas Butz. "Cardiogram: Visual Analytics for Automotive Engineers." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2011, pp. 1727–1736.

[242] Michael Sedlmair, Benjamin Kunze, Wolfgang Hintermaier, and Andreas Butz. "User-Centered Development of a Visual Exploration System for In-Car Communication." In: *Proceedings of the International Symposium on Smart Graphics*. 2009, pp. 105–116.

[243] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. "Design Study Methodology: Reflections from the Trenches and the Stacks." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2431–2440.

[244] Clifford A. Shaffer, Duane L. Knill, and Layne T. Watson. "Visualization for Multiparameter Aircraft Designs." In: *Proceedings of the IEEE Conference on Visualization*. 1998, pp. 491–494.

[245] Ben Shneiderman. "The Eyes Have it: A Task by Data Type Taxonomy for Information Visualizations." In: *Proceedings of the IEEE Symposium on Visual Languages*. 1996, pp. 336–343.

[246] Ben Shneiderman and Catherine Plaisant. "Strategies for Evaluating Information Visualization Tools: Multi-Dimensional In-Depth Long-Term Case Studies." In: *Proceedings of the AVI workshop on BEyond time and errors: novel evaluation methods for information visualization*. 2006, pp. 1–7.

[247] Siegfried Silber, Gerd Bramerdorfer, Alexander Dorninger, Armin Fohler, Johannes Gerstmayr, Werner Koppelstätter, Daniel Reischl, Günther Weidenholzer, and Simon Weitzhofer. "Coupled Optimization in MagOpt." In: *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering* 230.4 (2016), pp. 291–299.

[248] Siegfried Silber, Werner Koppelstätter, Günther Weidenholzer, Gordan Segon, and Gerd Bramerdorfer. "Reducing Development Time of Electric Machines with SyMSpace." In: *Proceedings of the International Electric Drives Production Conference*. 2018, pp. 1–5.

[249] Hillary Sillitto, Regina Griego, Eileen Arnold, Dov Dori, James Martin, Dorothy McKinney, Patrick Godfrey, Daniel Krob, and Scott Jackson. "What do we mean by "system"?–System Beliefs and Worldviews in the INCOSE Community." In: *Proceedings of the INCOSE International Symposium*. 2018, pp. 1190–1206.

[250] Herbert A. Simon. "A Behavioral Model of Rational Choice." In: *The Quarterly Journal of Economics* 69.1 (1955), pp. 99–118.

[251] Herbert A. Simon. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. Second. Macmillan, 1957.

[252] Herbert A. Simon. *The New Science of Management Decision*. Harper & Brothers, 1960.

[253] Herbert A. Simon. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. Third. Free Press, 1976.

[254] Johannes Sorger, Thomas Ortner, Christian Luksch, Michael Schwärzler, Eduard Gröller, and Harald Piringer. "LiteVis: Integrated Visualization for Simulation-Based Decision Support in Lighting Design." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2015), pp. 290–299.

[255] Bob Spence, Lisa Tweedie, Huw Dawkes, and Hua Su. "Visualisation for Functional Design." In: *Proceedings of the IEEE Symposium on Information Visualization*. 1995, pp. 4–10.

[256] Robert Spence. *Information Visualization*. 1st. ACM Press, 2001.

[257] Michael Spenke, Christian Beilken, and Thomas Berlage. "FOCUS: the Interactive Table for Product Comparison and Selection." In: *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*. 1996, pp. 41–50.

[258] Rainer Splechtna, Michael Beham, Denis Gračanin, María Luján Ganuza, Katja Bühler, Igor Sunday Pandžić, and Krešimir Matković. "Cross-Table Linking and Brushing: Interactive Visual Analysis of Multiple Tabular Data Sets." In: *The Visual Computer* 34.6 (2018), pp. 1087–1098.

[259] Rainer Splechtna, Krešimir Matković, Denis Gračanin, Mario Jelović, and Helwig Hauser. "Interactive Visual Steering of Hierarchical Simulation Ensembles." In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. 2015, pp. 89–96.

[260] A. Corbin Strauss and Juliet Corbin. *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage Publications, 1990.

[261] Uzma H. Syeda, Prasanth Murali, Lisa Roe, Becca Berkey, and Michelle A. Borkin. "Design Study "Lite" Methodology: Expediting Design Studies and Enabling the Synergy of Visualization Pedagogy and Social Good." In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–13.

[262] Sławomir K. Tadeja, Timoleon Kipouros, Yupu Lu, and Per O. Kristensson. "Supporting Decision Making in Engineering Design Using Parallel Coordinates in Virtual Reality." In: *AIAA Journal* 59.12 (2021), pp. 5332–5346.

[263] Yee S. Tan and Niall M. Fraser. "The Modified Star Graph and the Petal Diagram: Two New Visual Aids for Discrete Alternative Multicriteria Decision Making." In: *Journal of Multi-Criteria Decision Analysis* 7.1 (1998), pp. 20–33.

[264] Yuzuru Tanahashi and Kwan-Liu Ma. "Design Considerations for Optimizing Storyline Visualizations." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2679–2688.

[265] Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, 2008.

[266] Churee Theetranont, Peter Haddawy, and Donyaprueth Krairit. "Integrating Visualization and Multi-Attribute Utility Theory for Online Product Selection." In: *International Journal of Information Technology & Decision Making* 6.04 (2007), pp. 723–750.

[267] James J. Thomas and Kristin A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society Press, 2005.

[268] Thomas Torsney-Weir, Ahmed Saad, Torsten Möller, Hans-Christian Hege, Britta Weber, Jean-Marc Verbavatz, and Steven Bergner. "Tuner: Principled Parameter Finding for Image Segmentation Algorithms Using Visual Response Surface Exploration." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 1892–1901.

[269] Thomas Torsney-Weir, Michael Sedlmair, and Torsten Möller. "Decision Making in Uncertainty Visualization." In: *Proceedings of the IEEE Workshop on Visualization for Decision Making under Uncertainty*. 2015.

[270] Melanie Tory and Torsten Möller. "Human Factors in Visualization Research." In: *IEEE Transactions on Visualization and Computer Graphics* 10.1 (2004), pp. 72–84.

[271] Melanie Tory and Torsten Möller. "Evaluating Visualizations: Do Expert Reviews Work?" In: *IEEE Computer Graphics and Applications* 25.5 (2005), pp. 8–11.

[272] Melanie Tory and Sheryl Staub-French. "Qualitative Analysis of Visualization: A Building Design Field Study." In: *Proceedings of the Workshop on BEyond time and errors: novel evaLuation methods for Information Visualization*. 2008, pp. 1–8.

[273] Thomas Trautner and Stefan Bruckner. "Line Weaver: Importance-Driven Order Enhanced Rendering of Dense Line Charts." In: *Computer Graphics Forum* 40.3 (2021), pp. 399–410.

[274] Hans L. Trinkaus and Thomas Hanne. "knowCube: a Visual and Interactive Support for Multicriteria Decision Making." In: *Computers & Operations Research* 32.5 (2005), pp. 1289–1309.

[275] Alexis Tsoukiàs. "From Decision Theory to Decision Aiding Methodology." In: *European Journal of Operational Research* 187.1 (2008), pp. 138–161.

[276] Tea Tušar and Bogdan Filipič. "Visualization of Pareto Front Approximations in Evolutionary Multiobjective Optimization: A Critical Review and the Prosection Method." In: *IEEE Transactions on Evolutionary Computation* 19.2 (2014), pp. 225–245.

[277] Amos Tversky. "Elimination by Aspects: A Theory of Choice." In: *Psychological Review* 79.4 (1972), p. 281.

[278] Amos Tversky and Daniel Kahneman. “The Framing of Decisions and the Psychology of Choice.” In: *Science* 211.4481 (1981), pp. 453–458.

[279] Lisa Tweedie, Bob Spence, Huw Dawkes, and Hua Su. “The Influence Explorer.” In: *Proceedings of the Conference Companion on Human Factors in Computing Systems*. 1995, pp. 129–130.

[280] Lisa Tweedie, Bob Spence, David Williams, and Ravinder Bhogal. “The Attribute Explorer.” In: *Proceedings of the Conference Companion on Human Factors in Computing Systems*. 1994, pp. 435–436.

[281] David G. Ullman. “Robust Decision-Making for Engineering Design.” In: *Journal of Engineering Design* 12.1 (2001), pp. 3–13.

[282] Andrea Unger, Sven Schulte, Volker Klemann, and Doris Dransch. “A Visual Analysis Concept for the Validation of Geoscientific Simulation Models.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2216–2225.

[283] Eliane R. A. Valiati, Marcelo S. Pimenta, and Carla M. D. S. Freitas. “A Taxonomy of Tasks for Guiding the Evaluation of Multidimensional Visualizations.” In: *Proceedings of the AVI workshop on BEyond time and errors: novel evaluation methods for information visualization*. 2006, pp. 1–6.

[284] Emily Wall, Subhajit Das, Ravish Chawla, Bharath Kalidindi, Eli T. Brown, and Alex Endert. “Podium: Ranking Data Using Mixed-Initiative Visual Analytics.” In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2018), pp. 288–297.

[285] Junpeng Wang, Subhashis Hazarika, Cheng Li, and Han-Wei Shen. “Visualization and Visual Analysis of Ensemble Data: A Survey.” In: *IEEE Transactions on Visualization and Computer Graphics* 25.9 (2018), pp. 2853–2872.

[286] Junpeng Wang, Xiaotong Liu, Han-Wei Shen, and Guang Lin. “Multi-Resolution Climate Ensemble Parameter Analysis with Nested Parallel Coordinates Plots.” In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), pp. 81–90.

[287] Xiaoyu Wang, Dong H. Jeong, Wenwen Dou, Seok-Won Lee, William Ribarsky, and Remco Chang. “Defining and Applying Knowledge Conversion Processes to a Visual Analytics System.” In: *Computers & Graphics* 33.5 (2009), pp. 616–623.

[288] Matthew O. Ward. “Creating and Manipulating N-Dimensional Brushes.” In: *Proceedings of Joint Statistical Meeting*. 1997, pp. 6–14.

[289] Matthew O. Ward, Georges Grinstein, and Daniel Keim. *Interactive Data Visualization: Foundations, Techniques, and Applications*. Second. CRC Press, 2015.

[290] Colin Ware. *Information Visualization: Perception for Design*. Fourth. Elsevier, 2019.

[291] Jürgen Waser, Artem Konev, Bernhard Sadransky, Zsolt Horváth, Hrvoje Ribičić, Robert Carnecky, Patrick Kluding, and Benjamin Schindler. "Many Plans: Multidimensional Ensembles for Visual Decision Support in Flood Management." In: *Computer Graphics Forum* 33.3 (2014), pp. 281–290.

[292] Stephen Wehrend and Clayton Lewis. "A Problem-Oriented Classification of Visualization Techniques." In: *Proceedings of the IEEE Conference on Visualization*. 1990, pp. 139–143.

[293] Daniel K. I. Weidele. "Conditional Parallel Coordinates." In: *Proceedings of the IEEE Visualization Conference*. 2019, pp. 221–225.

[294] Di Weng, Chengbo Zheng, Zikun Deng, Mingze Ma, Jie Bao, Yu Zheng, Mingliang Xu, and Yingcai Wu. "Towards Better Bus Networks: A Visual Analytics Approach." In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), pp. 817–827.

[295] Di Weng, Heming Zhu, Jie Bao, Yu Zheng, and Yingcai Wu. "HomeFinder Revisited: Finding Ideal Homes with Reachability-Centric Multi-Criteria Decision Making." In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–12.

[296] Jarke J. van Wijk. "Bridging the Gaps." In: *IEEE Computer Graphics and Applications* 26.6 (2006), pp. 6–9.

[297] Christopher Williamson and Ben Shneiderman. "The Dynamic HomeFinder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System." In: *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1992, pp. 338–346.

[298] Kent Wittenburg, Tom Lanning, Michael Heinrichs, and Michael Stanton. "Parallel Bargrams for Consumer-Based Information Exploration and Choice." In: *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*. 2001, pp. 51–60.

[299] Meng-Jia Wu, Kelly Zhao, and Francisca Fils-Aime. "Response Rates of Online Surveys in Published Research: A Meta-Analysis." In: *Computers in Human Behavior Reports* 7 (2022), p. 100206.

[300] Kai Xu, Alvitta Ottley, Conny Walchshofer, Marc Streit, Remco Chang, and John Wenskovitch. "Survey on the Analysis of User Interactions and Visualization Provenance." In: *Computer Graphics Forum* 39.3 (2020), pp. 757–783.

[301] Chi-Hsien Yen, Haocong Cheng, Grace Y.-C. Yen, Brian P. Bailey, and Yun Huang. "Narratives+ Diagrams: An Integrated Approach for Externalizing and Sharing People's Causal Beliefs." In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW2 (2021), pp. 1–27.

[302] Ji S. Yi. *Visualized Decision Making: Development and Applications of Information Visualization Techniques to Improve Decision Quality of Nursing Home Choice*. Georgia Institute of Technology, 2008.

[303] Ji S. Yi, Youn-ah Kang, John Stasko, and Julie A. Jacko. "Toward a Deeper Understanding of the Role of Interaction in Information Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1224–1231.

[304] Ji S. Yi, Rachel Melton, John Stasko, and Julie A. Jacko. "Dust & Magnet: Multivariate Information Visualization Using a Magnet Metaphor." In: *Information Visualization* 4.4 (2005), pp. 239–256.

[305] Xiaoru Yuan, Peihong Guo, He Xiao, Hong Zhou, and Huamin Qu. "Scattering Points in Parallel Coordinates." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1001–1008.

[306] Alexandru-Ciprian Zăvoianu, Gerd Bramerdorfer, Edwin Lughofer, Siegfried Silber, Wolfgang Amrhein, and Erich Peter Klement. "A Hybrid Soft Computing Approach for Optimizing Design Parameters of Electrical Drives." In: *Proceedings of the International Conference on Soft Computing Models in Industrial and Environmental Applications*. 2013, pp. 347–358.

[307] Klaus Zeman et al. "Symbiotic Mechatronics: An Alternative View On Complex Systems." In: *Proceedings of the International Tools and Methods of Competitive Engineering Symposium*. 2020, pp. 11–15.

[308] Jian Zhao, Michael Glueck, Petra Isenberg, Fanny Chevalier, and Azam Khan. "Supporting Handoff in Asynchronous Collaborative Sensemaking Using Knowledge-Transfer Graphs." In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2017), pp. 340–350.

[309] Xun Zhao, Yanhong Wu, Weiwei Cui, Xinnan Du, Yuan Chen, Yong Wang, Dik Lun Lee, and Huamin Qu. "SkyLens: Visual Analysis of Skyline on Multi-Dimensional Data." In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2017), pp. 246–255.

[310] Hong Zhou, Xiaoru Yuan, Huamin Qu, Weiwei Cui, and Baoquan Chen. "Visual Clustering in Parallel Coordinates." In: *Computer Graphics Forum* 27.3 (2008), pp. 1047–1054.

[311] Michelle X. Zhou and Steven K. Feiner. "Visual Task Characterization for Automated Visual Discourse Synthesis." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1998, pp. 392–399.

[312] CORE studio. *Design Explorer (v2)*. http://tt-acm.github.io/DesignExplorer/. [Accessed 10 November 2023]. 2019.