

A Typology of Decision-Making Tasks for Visualization

Camelia D. Brumar ¹, Sam Molnar ², Gabriel Appleby ², Kristi Potter ², *Member, IEEE*,
and Remco Chang ¹, *Member, IEEE*

Abstract—Despite decision-making being a vital goal of data visualization, little work has been done to differentiate decision-making tasks within the field. While visualization task taxonomies and typologies exist, they often focus on more granular analytical tasks that are too low-level to describe large complex decisions, which can make it difficult to reason about and design decision-support tools. In this paper, we contribute a typology of decision-making tasks that were iteratively refined from a list of design goals distilled from a literature review. Our typology is concise and consists of only three tasks: CHOOSE, ACTIVATE, and CREATE. Although decision types originating in other disciplines exist, we provide definitions for these tasks that are suitable for the visualization community. Our proposed typology offers two benefits. First, the ability to compose and hierarchically organize the tasks enables flexible and clear descriptions of decisions with varying levels of complexities. Second, the typology encourages productive discourse between visualization designers and domain experts by abstracting the intricacies of data, thereby promoting clarity and rigorous analysis of decision-making processes. We demonstrate the benefits of our typology through four case studies, and present an evaluation of the typology from semi-structured interviews with experienced members of the visualization community who have contributed to developing or publishing decision support systems for domain experts. Our interviewees used our typology to delineate the decision-making processes supported by their systems, demonstrating its descriptive capacity and effectiveness. Finally, we present preliminary findings on the usefulness of our typology for visualization design.

Index Terms—Decision-making, typology, theory.

I. INTRODUCTION

VISUALIZATION researchers and practitioners alike have suggested that one of the primary goals of visualization systems is to help users make decisions [1], [2], [3], [4]. Many works have come out of the visualization community that have succeeded in helping users make important decisions (e.g., [5],

Received 9 September 2024; revised 5 May 2025; accepted 16 May 2025. Date of publication 22 May 2025; date of current version 5 September 2025. This work was supported in part by the National Renewable Energy Laboratory for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Recommended for acceptance by J. Zhao.

Camelia D. Brumar and Remco Chang are with Tufts University, Medford, MA 02155 USA (e-mail: Camelia_Daniela.Brumar@tufts.edu; Remco@cs.tufts.edu).

Sam Molnar, Gabriel Appleby, and Kristi Potter are with National Renewable Energy Laboratory, Golden, CO 80401 USA (e-mail: Sam.Molnar@nrel.gov; GAppleby@nrel.gov; Kristi.Potter@nrel.gov).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TVCG.2025.3572842>, provided by the authors.

Digital Object Identifier 10.1109/TVCG.2025.3572842

1077-2626 © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

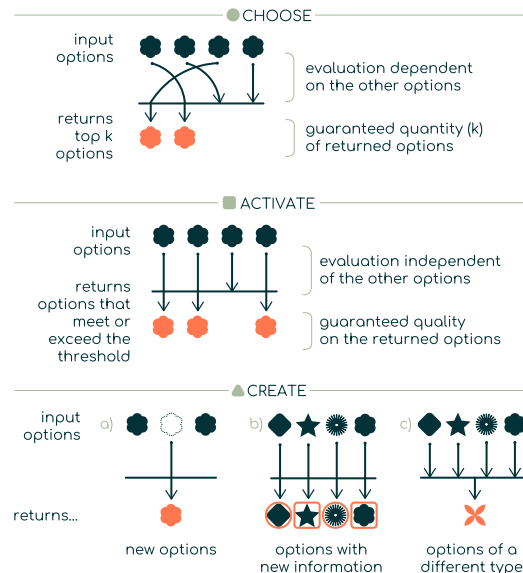


Fig. 1. The three decision tasks: CHOOSE, ACTIVATE, and CREATE.

[6], [7], [8], [9], [10], to name a few). However, it is difficult to understand patterns across these systems, even among those that address similar decision-making problems in the same domain. Furthermore, the taxonomies within the visualization community (e.g., [4], [11], [12], [13], [14], [15]) do not adequately capture the components of complex decision-making problems and processes because they instead focus on analytical tasks or visual encodings, or they categorize all decision-making tasks under a single “decide” or “choice” task. This leaves a critical gap in understanding how information is processed when making a decision and how individual decisions combine to form a coherent decision-making process.

To address these shortcomings, we introduce a typology specifically focused on decision-making tasks in visualization. It consists of three decision tasks: CHOOSE, ACTIVATE, and CREATE, as shown in Fig. 1. Each task assesses options and represents distinct decisions with those options. The CHOOSE task decides which options are optimal or best, the ACTIVATE task decides which options meet or exceed a threshold, and the CREATE task decides how to assemble, synthesize, or generate new information from the options. The most critical properties of this typology are that it allows us to represent the complex structures of real-world decision-making problems through

the hierarchical organization of tasks and the composition of tasks.

We designed our typology using design goals derived from a literature review of decision-support tools within the visualization community. We built upon previously curated corpora of application and decision-support visualization systems [15], [16], focusing on those with domain expert evaluations. We then followed an inductive coding process to identify each paper's primary decision-making goals and processes. From these codes, we distilled the design goals for our decision-making typology, including iterativeness, composability, generalizability, simplicity, and others. From decision-making theory [17], [18] and our design goals, we worked towards an initial set of tasks for the typology. We used our survey to hone and iterate on our tasks until they could accurately characterize the works included. To illustrate that our typology has the aforementioned properties, we present four case studies where authors described the decision-making processes addressed by their tools using our typology.

We validated our typology by conducting semi-structured interviews. Our nine participants were visualization experts from academia, government laboratories, and industry who had authored one or more decision-support tools. We demonstrated our typology to the participants and asked them to apply it to their systems and answer open-ended questions about the typology, its properties, and how it compares to other typologies.

We divide the results of our validation into three broad categories: the correctness of the typology, the utility of the typology, and how the typology compares to existing task taxonomies and typologies. The correctness of the typology can be further broken down into the consideration of completeness (the typology needs to cover all decision-making problems we have encountered) and expressiveness (the typology can be easily used to describe decision-making scenarios). In our interview study, participants found that the typology is complete regarding the three decision tasks and that they could successfully express their decision-making problems in diagrams using the typology. They also found the typology helpful when analyzing and reflecting on the design process of visualization decision-support tools, believed that it could be useful when communicating with domain experts, and noted that it could categorize visualization decision-making publications. Moreover, participants provided valuable insights when comparing the typology to existing frameworks, highlighting its appropriate level of abstraction for describing decision-making problems and its adaptability to diverse data and contexts.

Lastly, we conducted a preliminary study to evaluate the impact of our typology on visualization design. Our findings revealed that participants used the typology to guide their choices between design alternatives, identify improvements for decision tasks that lacked sufficient support, and even redesign their systems based on the insights gained from applying the typology. With the results from our evaluation, we gathered all the insights as limitations and possible future avenues for our typology.

In summary, the major contributions of this paper are:

- Design goals for decision-making typologies derived from the thematic analysis of a comprehensive literature review.
- A typology of decision tasks, which provides an abstract yet precise way to discuss decision-making activities in visualization.
- Four case studies in which our typology is used to describe users' decisions in the context of decision-support tools that demonstrate the properties of our typology.
- A semi-structured interview study that serves as an evaluation of our typology and provides a roadmap for future work.
- A preliminary study on the design implications of our typology and discussion on the utility of this typology in the visualization community.

II. BACKGROUND & RELATED WORK

Decision making is a broad topic studied across many disciplines, and visualization researchers emphasize that supporting decision making is a vital goal of data visualization [4], [19], [20], [21], [22]. To develop effective support, the visualization community utilizes typologies as tools to reason about visualization designs and explain problems and user tasks addressed by visualization. We review the relevant literature from these areas and discuss the significance to this work.

A. Decision Making & Visualization

The visualization community tends to focus on decision making from three broad lenses: cognition and perception, decision process support, and domain-specific decision support systems. In a cognitive context, decision-making is often framed in terms of rationality and judgement [23]. In this view, the most rational decisions are ones that maximize expectation of a function [24]. However, a number of human decisions are not rational [25], a fact that has impacted the research proclivities of the visualization community in two significant ways. The first is an increase in research on computational decision-making support [19], [26], [27]. While promising, it is unclear and often difficult to determine under what conditions users will accept help from or implement decisions of computational models [10], [28]. The second impact is the increased study of cognitive bias and perception in visualization contexts. Cognitive biases explain irrational decision-making as the result of using heuristic functions to evaluate information in a decision-making cognitive process [29].

The occurrence of specific kinds of judgment and decision-making errors can then be traced to specific heuristic functions, and recent studies have investigated the overlap of specific biases with visualization design [30]. There is evidence that visualization design can impact the strength of a cognitive bias [31], and there exists a significant body of work that assesses the role of visual perception in judgment and decision-making [23], [32], [33]. For example, in Padilla et al. [34], the authors find a significant relationship between participants' risk assessment and uncertainty visualization design choices. While important, cognitive and perceptual approaches focus on reasoning and

judgement, rather than describing decision-making problems in domain agnostic contexts for meaningful discussions on visual support.

A second body of work in the visualization community is concerned with designing visualizations that support specific problem-solving structures for a decision-making problem. These visualization systems are unique in that they support a specific approach for decision-making rather than focusing on a specific application or domain [15], [35], [36], [37], [38]. One such decision process addressed in the community is multiple criteria decision making (MCDM) which is a mathematical formalism that enables selection of a solution given multiple competing criteria and a decision maker's preferences for various attributes of a solution [39]. MCDM lends itself nicely to computation and is particularly useful for ranking solutions [40], defining (and therefore exploring) trade-offs among solutions [35], and the explicit use of a person's preferences can be an invaluable source of information for recommendation systems or group decision-making [41]. However, a prescribed structure to a decision-making problem is not always appropriate, and transforming a decision-making problem to fit into a particular structure can itself be challenging for users. For example, MCDM hinges on the specification of criteria and evaluation functions for those criteria [39]. However, for some problems the selection of criteria and evaluation function is in itself difficult [35], [42], [43], and externalizing preferences as abstract values can be a hard interpretation problem for users [44]. Focusing solely on describing and discussing problems from a specific structure unnecessarily constrains possible visualization solutions, and does not capture the diverse ways people can approach complex decision-making problems.

Finally, there exists a large body of work in the visual analytics community which addresses decision-making from domain-specific contexts. For example, many decision support systems exist to address medical decisions [8], [45], [46], security and fraud detection [7], [47], [48], [49], city planning [9], [50], [51], and other domain-specific applications [52], [53], [54], [55], [56]. When supporting decision making within a particular domain, it is indeed important to take into account the nuances relevant to that domain. However, domain-specific approaches do not necessarily generalize, and do not allow for a discussion of decision-making problems across domains.

B. Visualization Typologies

Decision-making happens at many different levels of interaction with data and within visualization systems. Yet, no existing visualization typology can capture both high-level decisions (e.g., what car to buy) and low-level decisions (e.g., where do I click next in the application to show cars with specific attributes?) while simultaneously representing relevant decisions that occur outside of those systems. We briefly review prominent and relevant typologies of the visualization community here and discuss why these typologies are insufficient for discussing decision-making.

Two of the most prominent and influential visualization task typologies in the community are the analytical task typology by

Amar et al. [11] and the multi-level task typologies by Brehmer and Munzner [13], [57]. The analytical tasks typology addressed a gap in highly data-centric typologies by taking an analytical-centric approach. This typology consists of ten tasks that describe "primitive" *analyses of data*, or low-level tasks [11]. The multi-level task typology further expanded upon these tasks to incorporate higher levels like "why?", "how?", and "what?", which can be composed from lower-level tasks [13], [57]. However, these tasks are not suitable for describing decision-making in compact yet meaningful ways. Consider the decision of buying a car by choosing among a set of available used cars. This task, in its current form, cannot yet be represented in either the analytical or multi-level task typologies because more information is needed on the analytical activities and data. In other words, these typologies are oriented towards describing potential *solutions* or approaches for solving a problem, they do not describe the decision *problem* or how decisions are made. When we can describe decision problems in abstract yet precise ways, we can begin to reason about solutions more holistically.

III. METHODOLOGY

This section describes our methodology for developing a typology of decision-making tasks. Our approach involved gathering design goals from a literature review of publications in the following fields: visualization relating to decision-support tools, cognitive psychology relating to mental processes for decision-making, and other scientific literature relating to decision-making. Specifically, we iteratively analyzed design studies and visual analytics application papers to extract example scenarios and relevant contexts relating to decision-making, as detailed in Section III-A and III-B. This analysis led to a set of design goals we believe our typology should fulfill in order to accurately represent the decision processes described in the corpus's case studies and decision-making scenarios. These design goals are described in Section III-C. Next, we shaped the typology using theory frameworks for decision-making by keeping in mind the goals we derived from the application papers, as explained in Section III-D. The result of this design process is our proposed typology of decision-making tasks for visualization, which we present in detail in Section IV.

A. Application Survey Corpus

We reviewed relevant visualization decision-support literature to formulate the design goals of our typology, building upon the literature corpora curated by Oral et al. [15] and Shi et al. [16]. Oral et al.'s corpus focused on papers explicitly labeled as decision-making tools, filtering based on the presence of the term 'decision' in titles or abstracts. However, this approach omitted valuable visualization papers which serve as decision-support tools for domain experts and tools implementing methodologies other than MCDM. To broaden the scope, we incorporated Shi et al.'s work, allowing us to encompass a more comprehensive range of visual analytics applications and design studies. From their list of 190 VAST papers [16], we found 63 to be applications with a direct decision-making focus. We added to that list the

papers with “decision” mentioned in titles or abstracts to be consistent with [15], adding up to 74 papers.

To refine our selection, we prioritized papers featuring specific domain applications, and eliminated generic systems lacking evaluations with domain experts. This guaranteed that the chosen tools were utilized by domain users with actual decisions and requirements that could be validated. For instance, papers solely based on hypothetical case studies were excluded (e.g., Podium [37], Zooids [58]). This filtering process resulted in a final corpus of 69 papers that directly contribute to our study’s objectives, and is available in the Supplemental Materials.

B. Coding Process

Our coding process followed an inductive approach to derive the design goals for our typology from the survey corpus outlined in the preceding section. The process involved a team-based approach with two coders, who are the first two authors of this paper. For each paper, they determined the overall decision goal by analyzing the task abstraction, requirements, and evaluation sections, as well as any supplementary materials, such as figures or video demonstrations (if available). The identified decision goals of the tools informed the design goals for a task typology. In addition, we noted unique decision contexts, such as sub-decisions that supported the overall decision, and the information that decisions use and produce. The result of this coding process resulted in seven design goals for our typology, which are discussed in the next subsection.

C. Design Goals for a Decision Task Typology

Through multiple brainstorming sessions involving all authors, we distilled the following list of goals that we believe a decision-making typology for visualization should fulfill in order to accurately represent the decision processes and scenarios found in the corpus. In the supplemental materials, we include excerpts from the corpus of papers that support the choice of each goal.

- G1 Iteration:* Given that decision-making processes in visualization often involve cycles and iterations, our typology should incorporate a mechanism to express such iterative patterns.
- G2 Hierarchical Structure:* Decisions are intricate and can be deconstructed into subtasks, necessitating the ability of our typology to capture these hierarchical relationships.
- G3 Process-Centric Approach:* Our typology should focus on delineating the sequence of decisions users encounter, rather than merely representing data-processing flow or interaction order typically outlined in task lists and design requirements of application papers and design studies.
- G4 Information Processing-Focused:*

The decision tasks within our typology should demonstrate flexibility and impartiality towards input and output types, instead emphasizing how information is processed and evaluated.

Additionally, we define the following evaluation criteria, which we use to assess the effectiveness of our typology through user evaluations:

E1 Generalizability: The typology should apply uniformly across the diverse range of papers we surveyed and should remain relevant across various domains.

E2 Simplicity: We aspire for our typology to remain straightforward and concise, with minimal decision types and operations to avoid unnecessary complexity.

E3 Accessibility: Our typology should be easily understood and grasped by both the visualization designers and the domain experts, contributing to its usability and adoption.

These criteria guided the creation of our typology, as we discuss in the following subsection.

D. Designing the Typology

Our design process began with an initial set of decision-making tasks inspired by relevant decision-making theory literature, particularly works by Fischhoff et al. [17], [24] and Simon’s decision-making phases [18], [59]. Then, using the goals from the previous subsection and other theory found in cognitive science literature [23], [25], [29], [60], we iteratively added, removed and modified the decision tasks and their properties until we achieved a stable set of decision-making tasks and definitions, resulting in the finalized typology described in Section IV.

In terms of the different types of decision-making tasks, previous works often describe a single decision task, *decide* [4], [29] or *choice* [15], [60], [61], both referring to the action or process of deciding or resolving something. While useful for encapsulating decisions as a whole, they lack descriptiveness regarding the type of decision, that is, *how* a decision is made, limiting the analysis of decision-making in visualization [29].

There are two prominent works that are relevant to consider which outline possible decision tasks. We tested each of the decision tasks outlined in these works by applying them to our corpus. For example, we conducted an initial pass over the decision-making support tools and labeled them with one of the decision tasks, attempting to represent the ultimate decision the tool supports. For each categorization, we noted which of our design goals they did and did not meet.

First, Fischhoff et al. [17] describe three distinct decisions in the form of questions: “Is it time to act?”, “Which is best?”, and “What is possible?”. While these better describe types of decisions in a general and domain-agnostic way (**E1**), the associated definitions are not directly applicable to visualization because they lack specificity on the use of information to reach conclusions or achieve decision goals (**G4**), which we found distinguished decisions and visualization designs in our corpus. Additionally, as defined by Fischhoff et al. these tasks are not composable (**G1 - G3**), which limits our ability to describe complex decision-making processes in sufficient detail for designing decision-support systems (**E1 - E3**).

The second prominent work outlining decision types is by Herbert Simon, which describes decision-making as a process comprising three stages: *intelligence*, *design*, and *choice*. The *intelligence* stage involves gathering information and identifying problems or opportunities that require a decision, the *design*

stage involves developing, analyzing, and evaluating options, and the *choice* stage involves selecting and implementing one of the options, considering constraints and objectives. According to Simon, each of these stages can be further broken down into sub-stages of the same type. While Simon’s model meets goals **G1** - **G2** and has property **E1**, it does not specify how information is processed within these stages (**G4**) or how the stages interrelate to achieve the final decision (**G3**).

Fischhoff and Simon’s decision types meet some of our typology design goals, but both fail to delineate differences in information processing for decisions, and neither discuss the composition of decisions in sufficient detail for visualization. Knowing the specifics of how information is processed, and therefore how decisions are made, and how smaller decisions compose together to reach a final decision, allows designers to create visualizations that better support those processes. We expand on past work by explicitly addressing how information is processed in our tasks (**G1-4**), and ensure the aforementioned properties exist in our typology, which we discuss next.

IV. TYPOLOGY OF DECISION-MAKING TASKS

Our typology consists of three tasks, which we derived from examining the scientific literature and iteratively refining the definitions to meet our design goals (see Section III). These tasks reflect and capture the ubiquity of decision-making, flexibly describe the execution of decisions, and the definitions use language consistent with the decision-making literature. The tasks are agnostic to who or what performs the evaluation, which ensures the typology can describe human or computational decision-making problems. Furthermore, our task definitions are *agnostic to the rationality* of the evaluations, which allows for our typology to be consistent and used in conjunction with computational and cognitive biases [29], [62]. We illustrate each decision task visually in Fig. 1.

A. Task Definitions

Each task in our typology evaluates and returns options according to its definition. *Options* are information entities typically relevant for a decision task. Options can have *features* that describe characteristics of the options. Lastly, *criteria* are preferences or standards on features of the options. As an example, we could consider options as cars, with features of color, price, and mileage, and an example criteria is “any color except white”.

1) **CHOOSE**: The CHOOSE task involves assessing a set of options and returning a subset deemed optimal or best. Fischhoff’s concept of *decisions with fixed options*, and Simon’s final decision-making stage, *choice*, form the basis of the CHOOSE task. We more specifically define CHOOSE as deciding on the top or best k options available (illustrated at the top of Fig. 1). Importantly, within CHOOSE, the evaluation of options is *dependent* on the other options. Whether an option is returned depends on how it compares with the other available options.

Consider the example decision of buying a car; we have four cars from which we can CHOOSE. Rather than looking for a car that satisfies a set of criteria, we are looking for the best

car *from the current set available*. Furthermore, the notion of *best* can change over time, which can lead to different choices. These changes could be due to different environmental or circumstantial contexts (*e.g.*, living in a new location), or shifting perceptions (*e.g.*, brand loyalty). Additionally, cognitive biases like loss aversion or the framing effect [29], and machine biases like the linking bias [62], are examples of biased evaluations that can occur in the CHOOSE task.

Unlike the other tasks in our typology, with CHOOSE we know how many options are returned upon execution, meaning only one execution of the CHOOSE task is required to reach a subset of options of size k . Although criteria can inform comparison of options in CHOOSE, criteria are not necessary to execute CHOOSE. Furthermore, it is not a given that the returned options meet a set of criteria. In other words, we have *no guarantees on the quality* of the returned options, only that they are the best subset of all evaluated options.

2) **ACTIVATE**: The ACTIVATE task draws inspiration from Fischhoff’s decisions about action thresholds but expands beyond merely temporal thresholds. This task represents a decision where options are evaluated, and only those that meet or exceed a threshold are returned (illustrated in the center of Fig. 1). In contrast to CHOOSE, the evaluation of one option is *independent* of the others. Whether an option is returned depends only on its evaluation individually, regardless of the evaluation of the other options.

In the car-buying example, the ACTIVATE task can be represented by applying thresholds to each car such that only cars under a particular price and mileage are left after evaluation. Notice that because each car is evaluated independently, we can assess them all in the evaluation *simultaneously*. There is flexibility in the evaluation that occurs in ACTIVATE, as thresholds need not be strictly enforced across options or well-defined. For example, in the car buying decision, a person may exclude a car that is very near (but still under) *both* the price and mileage limits, but keep a car that is slightly over the price limit because the mileage is much lower than the limit. Some example biases that can occur in this evaluation include the cognitive biases of information bias and the anchoring effect [29], or the algorithmic processing bias in computational decision-making [63].

Unlike CHOOSE, where we know the exact number of options that are returned, we do not have this guarantee for ACTIVATE. It is possible that none or all options are returned by ACTIVATE. We have to execute ACTIVATE an unknown number of times to reach a subset of options of size k . However, any options returned by ACTIVATE meet a level of acceptability that is informed by given criteria. In other words, ACTIVATE can *guarantee a level of quality* on the returned options that is not possible to do with CHOOSE.

3) **CREATE**: The CREATE task represents decisions on assembling, synthesizing, or generating new information. This decision represents the creation of knowledge or insights [1], [64], [65], [66], or more tangibly, reports [67], new data transformations [68], etc. In contrast with the other tasks in our typology, CREATE can take any number of options (including none) and return new options (*e.g.*, adding to the set of options or new features to existing options), options with new information

or modifications, or returning a set of options of a completely different type. These facets of the CREATE task are illustrated as a, b, and c, respectively, at the bottom of Fig. 1.

Consider again the decision to buy a car. Rather than deciding on an existing set of cars, the CREATE decision allows for three possibilities: (1) the addition of a new car that is not in the current set of options, (2) the modification or inclusion of features such as adding whether each car has advanced driver assist technologies, or (3) synthesizing new information based on the options, like creating a custom car where the initial set of options are components of a car (e.g., body, wheels, software capabilities). The last possibility is perhaps the most impactful, as it describes the case where the evaluation in CREATE returns a new type of option than the initial set of options.

To further illustrate (3), consider an evacuation planning scenario faced by emergency responders during a wildfire. One task involves using predicted fire spread (the input information) to determine which households should receive evacuation orders. While this could be framed as a task involving refining a list of all households (i.e., a CHOOSE or an ACTIVATE task), the decision can also be reframed as a CREATE task: evaluating one type of information (e.g., spatiotemporal fire forecasts) to generate a new type of actionable output (e.g., a list of households to evacuate). This example of transformation, from predictive modeling to actionable directives, highlights another core capability of the CREATE task: the output is not necessarily a refined version of the input but may be a fundamentally different representation, tailored to a different decision context. As discussed further in Section IV-B, the ability to reframe a decision as different decision-making tasks is a strength of the typology, as it supports reflection on the various ways information can be processed.

The CREATE task is often considered a precursor to decisions in the literature. In Simon's framework, creating alternative courses of action by evaluating their consequences falls under the *design* stage [18], [69]. However, in agreement with Fischhoff et al. [17], we argue that the CREATE task is more than just a precursor to a decision—it is itself a fundamental decision that can be the ultimate goal of a decision support tool. For instance, generating a report that synthesizes insights while using a tool exemplifies how the CREATE task serves as an end decision goal.

While we can make some guarantees on the quantity or quality of returned options in CHOOSE and ACTIVATE, we cannot do so for CREATE. If the initial options do not meet a set of criteria, this can *propagate* through the CREATE evaluation and impact the *quality* of the returned options. Furthermore, even if the initial options do meet a set of criteria, the returned options may require a different set of criteria in order to assess their quality. Constraints on the *quantity* of returned options from the CREATE evaluation could arise from initial options that do not meet a set of criteria. However, the occurrence of inspiration, innovation, or randomness in the CREATE evaluation can allow an unbounded number of returned options. The quality and quantity assertions for CHOOSE and ACTIVATE are possible because we can reason directly on the initial options in relation to their evaluations. While we can also reason about the impact of the initial options on the CREATE task, the CREATE evaluation is less bounded in that there is no prescribed way to use the initial options. This leads to additional uncertainty in the quality and

quantity of returned options than what is present in CHOOSE and ACTIVATE.

B. Discussion of Definitions

The three decision tasks are distinguished by how they evaluate information, or *how* a decision is made, thus meeting goal G4. The same decision *goal* (e.g., buying a car) is achievable with any decision *task* if the appropriate information (e.g., the cars, their features, and the criteria) is available and the evaluation is possible. For example, a user could prefer an ACTIVATE task because evaluating options individually by their features is easier than assessing options against one another. Alternatively, another user may prefer to execute a CHOOSE task because assessing an option in the context of the available options is helpful. Or perhaps yet another user prefers to execute a CREATE task because they prefer to devise their own option while drawing inspiration from the available ones. Although we use clear language distinguishing options from features and criteria, the decision tasks can operate on any of these entities. For example, as part of the decision goal of buying a car, we could CREATE the criteria that are then used to evaluate cars. The decision tasks allow us to describe how different users or algorithms approach decision-making, as well as how their strategies may shift in response to evaluation challenges or changes in information availability.

The decision tasks in our typology are independent from but synergistic with the intent (task framing), judgement (valuation of information), or analytical actions that can occur while executing the task evaluation. Specifically, these facets of decision-making do not change the *evaluation mechanics* of the decision tasks, but they can impact the *outcome* of a decision task. For example, an ACTIVATE task does not depend on whether a user executed it with the intent to include the best options or to exclude the worst options, as either intent can be achieved with an ACTIVATE evaluation. However, the intent may influence how each option is assessed and therefore the outcome of the ACTIVATE task. This distinction is further highlighted when considering the judgment and analysis of information. An anchoring bias in a CHOOSE task influences how later options are evaluated, but it does not change the details of the evaluation. Whether we use sorting or correlation as an analytical technique in CHOOSE only impacts which options are returned. Our decision tasks describe the unique functions of how decisions are made, while these other facets of decision-making help explain potential differences in decision outcomes.

The distinction on how information is processed to make a decision, rather than what influences the outcome of a decision, helps ensure our typology is simple, generalizable, and accessible (properties E1 - E3). Given our typology has the flexibility to describe a decision goal using any decision task, or combination of decisions tasks, hints at other important properties of our typology that help us meet goals G1 - G3, namely the composability and hierarchical structure of the tasks.

V. USING THE TYPOLOGY FOR DECISION-MAKING PROBLEMS

In isolation, each decision task already provides new ways to express decisions in relevant ways for visualization.

However, many decision problems are complex and involve multiple related decisions [69]. While a list of sub-decisions can partially describe larger decisions, it cannot capture the structural relationship that clearly exists among decisions. In order to meet goals **G1** - **G3**, we determined how our decision tasks can represent more complex and interrelated decisions. Specifically, the decision tasks are *functional descriptions* of information evaluation, which exposes two properties of our typology that enable us to combine tasks and represent decision-making problems: composability and divisibility. Furthermore, we can use a diagram to illustrate how these properties can describe more complex decision problems.

A. Properties That Help Describe Decision Problems

Consider again the example decision of wildfire evacuation: CHOOSE an evacuation location for residents. We could decompose this CHOOSE task into several sub-decision tasks like a CREATE task that decides a possible set of locations given predictions of where the fire may spread, and an ACTIVATE task that decides which locations are far enough from the fire zone yet still in the vicinity for evacuated residents, and enough capacity for everyone. The ability to break down the large CHOOSE task into sub-tasks illustrates how smaller decisions help resolve the overall decision. This example illustrates two properties of our typology:

Composability: The sub-decision tasks are interconnected and describe a flow among decisions.

For example, CREATE decides potential options and ACTIVATE then decides which options meet given criteria. This flow reflects how decision-making processes are structured as sequences of information transformations, supporting a process-centric view of decision-making (design goal **G3**). Moreover, the ordering of these tasks implies a sequence that captures iterative and dependent structures common in real-world decision-making (design goal **G1**). In the example above, the execution of CREATE must occur before the ACTIVATE decision because the information output by CREATE is the information evaluated in ACTIVATE. The property of composability is important for describing visualization tasks in general, as demonstrated by the task typology by Brehmer et al. [13]. Composability allows us to describe the relationship between decisions, which in turn helps us understand how information informs decisions (*e.g.*, we know the set of options output by the ACTIVATE decision are a good set to CHOOSE from) and is transformed by decisions (*e.g.*, predictions of wildfire spread into potential evacuation locations).

Decomposability: The decomposition of the CHOOSE task in the above example into sub-tasks implicitly defines a hierarchical structure of decision tasks (design goal **G2**). They provide a high-level abstraction of a decision task, while detailing the step-by-step decision-making processes at lower levels. The hierarchies indicate lower-level decisions that are needed to execute higher-level decisions due to its complexity. By expanding complex decisions into less complex constituent decisions, we can describe situations when multiple disparate decisions make-up a larger more complex decision.



Fig. 2. A diagram of the example wildfire evacuation decision-making problem. The CREATE task has an arrow into the ACTIVATE task, indicating that the CREATE task must occur before the ACTIVATE task. The ACTIVATE task has a self-loop, which indicates the possibility that this task is executed more than once. Lastly, the ACTIVATE task has an arrow into the CHOOSE task, indicating that the locations decided on in the ACTIVATE task are the input to the CHOOSE decision.

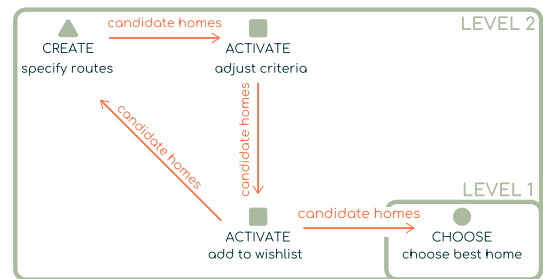


Fig. 3. Homefinder Revisited [70] decision diagram.

B. Diagramming Decision-Making Problems

We use a directed graph, referred to as a “diagram,” to illustrate decision problems represented in our typology. An example for the wildfire evacuation decision is shown in Fig. 2. In this diagram, decision tasks are depicted using distinct shapes: a circle for CHOOSE, a square for ACTIVATE, and a triangle for CREATE. These symbols are supplemented with free-form annotations in domain-specific language to provide details about the decisions being made. Arrows indicate the flow of information between decisions, with annotations on the edges specifying what information is being transferred. Additionally, the “LEVEL X” notation is employed within the case studies to represent the hierarchy of decisions: a LEVEL 1 decision corresponds to the overall goal, a LEVEL 2 decision supports the LEVEL 1 goal, a LEVEL 3 decision supports a LEVEL 2 decision, and so on.

VI. CASE STUDIES

To illustrate further our typology and its properties, we present four case studies where our typology is used to describe user decisions in the context of decision-support tools.

A. Case Study 1: Homefinder Revisited

The paper “Homefinder Revisited” [70] introduces ReACH (Reachability-Aided Contemporary HomeFinder), a visual analytics system aimed at aiding users in discovering their ideal homes. In ReACH, users input their daily routes, activities, and preferences, upon which the system generates potential homes by considering factors like reachability, proximity to amenities, and user-specified preferences. With its interactive interface, users can query, filter, and evaluate home candidates based on multiple criteria, ultimately choosing their ideal home. Fig.

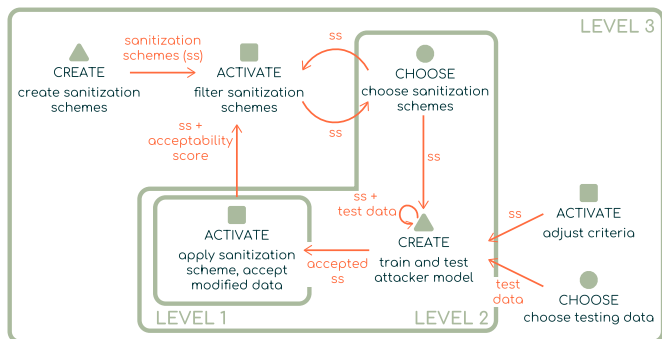


Fig. 4. Umbra [47] decision diagram.

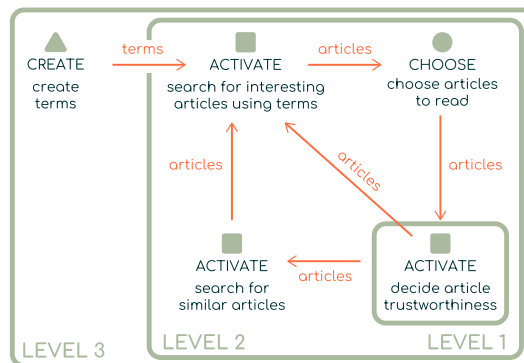


Fig. 5. Centaurus [71] decision diagram.

3 illustrates the following breakdown of the decision-making tasks:

Level 1: At the root level, ReACH supports the user to CHOOSE their ideal home from a number of candidates.

Level 2: Expanding on the CHOOSE decision, a user needs to first CREATE home candidates from their routes, activities, and preferences. The user performs the ACTIVATE decision to add or filter the candidates by refining the search criteria (e.g., parameters for reachability, preferences, etc.). The user then decides the homes that they would like to add to their wishlist (an ACTIVATE decision) and iterates on the process until the wishlist is complete. Finally, the user CHOOSES the ideal home from the wishlist.

Discussion: The decision-making framework described in the Homefinder Revisited paper represents a classic multi-criteria decision-making (MCDM) task. Typically, an MCDM task follows a diagram of CREATE (creating design criteria) → ACTIVATE (filtering and refining the criteria) → CHOOSE (selecting the best options). ReACH adds an additional decision task to this process, allowing users to add or remove candidates from a wishlist.

B. Case Study 2: Umbra

Umbra [47] is a visual analytics tool that helps address the complex problem of sanitizing sensitive information within datasets without compromising data characteristics, particularly in scenarios involving medical data, personally identifiable information, and classified data. Umbra is semi-automated in that it can: (1) automatically identify potential risks, (2) suggest sanitization actions and schemes, and (3) perform simulated attacks. Using Umbra, an analyst would review the potential risks, run simulated attacks, adopt or develop a sanitization scheme, and finally execute the scheme on a given dataset and disseminate the result. The decision tasks supported in Umbra can be viewed hierarchically and are illustrated in Fig. 4:

Level 1: At the root level, Umbra supports the ACTIVATE decision by evaluating whether the modifications to a sanitized dataset are acceptable, based on its remaining utility and the level of privacy achieved.

Level 2: At the next level, the ACTIVATE decision can be expanded to include two additional decisions: a CHOOSE decision, where the users choose what sanitization scheme to apply

to the dataset and where to apply it, and a CREATE decision, where a simulated attack is generated. This forms a loop: (1) choose a sanitization scheme, (2) run the simulation, and (3) accept or reject the sanitized data.

Level 3: The simple loop can be decomposed to include more nuanced decisions in the process. The CHOOSE decision of a sanitization scheme can be expanded into its own decision-flow that includes creating a sanitization approach (a CREATE decision), and deciding whether to accept the approach (an ACTIVATE decision). Similarly, the CREATE decision from Level 2 on generating a simulated attack can be decomposed further into deciding on a plan of attack (an ACTIVATE decision) and deciding whether to use custom or default test data (a CHOOSE decision).

Discussion: In the first case study, the “Homefinder Revisited” paper describes the fundamental user task at the root level as a CHOOSE decision for choosing an ideal home from a finite list of candidates. However, in this second case study, the number of candidates is potentially infinite, as Umbra can generate as many sanitization schemes as required. As a result, the Umbra system is designed to assist users in examining and refining a scheme, culminating in the decision of whether the scheme can be ACTIVATED and applied to sanitize a dataset.

C. Case Study 3: Centaurus

The Centaurus tool was utilized in a case study to evaluate the truthfulness of claims by analyzing related news articles [71]. Users were presented with a claim, such as a tweet, accompanied by a mix of trustworthy and untrustworthy news articles. Through Centaurus, users engaged with these articles to extract relevant information and assess the trustworthiness of each article. The primary focus was on understanding how users utilized Centaurus to formulate their final evaluations, irrespective of their stance on the original claim’s truthfulness. Fig. 5 illustrates the following breakdown of the decision-making tasks:

Level 1: The core of the Centaurus decision problem is an ACTIVATE decision, wherein users evaluate the trustworthiness of articles.

Level 2: The Level 1 ACTIVATE decision is supported by additional decisions – an ACTIVATE to select relevant articles based on search terms, a CHOOSE to choose articles for deeper




Fig. 6. WireVis [7] decision diagram.

analysis, and another ACTIVATE to search for similar articles after a user has made a decision about a potential article to investigate further.

Level 3: The ACTIVATE decision for selecting articles based on search terms is expanded to include a CREATE node where a user can manually enter keywords as search terms.

Discussion: While the domains of Centaurus (determining article trustworthiness) and Umbra (choosing data sanitization) may seem different, they exhibit notable similarities in the decision-making processes. Fig. 5 shows the decision diagram of Centaurus that uses a similar layout as the Umbra diagram in Fig. 4. By comparing the two, one can observe that Centaurus features the same three tasks at the top row as Umbra (CREATE, ACTIVATE, CHOOSE), and the bottom row also shares similarities. This comparison highlights the utility of our typology in evaluating and contrasting the decision tasks of seemingly disparate decision-support visualization tools.

D. Case Study 4: WireVis

The WireVis paper introduces a visual analytics system designed for the analysis of large amounts of financial wire transactions, aiming to uncover suspicious or fraudulent activities such as money laundering [7], [72]. The use of the WireVis system entails two distinct yet interconnected decision tasks: (1) identify fraudulent transactions and their associated accounts, and (2) discover previously unknown strategies used to commit financial fraud. In both cases, analysts use a coordinated set of visualizations to examine financial transactions from multiple perspectives, including keywords, amounts, time, and account relations. The system can automatically generate hierarchical clusters and perform pattern matching, thereby assisting analysts in their investigations.

Focusing on the task of “discovering previously unknown strategies used to commit financial fraud,” the decision tasks can be represented as shown in Fig. 6:

Level 1: The core objective can be described as a CREATE decision where an analyst aims to synthesize and discover a previously unknown “financial fraud strategy” based on suspicious financial transactions and activities.

Level 2: In order to CREATE the strategy, an analyst first need to CHOOSE the most promising accounts to investigate. For these accounts, the analyst determines the set of their transactions that might exhibit an unusual pattern of behavior (an ACTIVATE decision). These transactions are examined based on a list of criteria, such as the amount, the keyword used, the date patterns, etc. (an ACTIVATE decision). If deemed suspicious, an

analyst might be able to deduce a previously unknown strategy of financial fraud (a CREATE decision).

Discussion: The WireVis case study diverges from previous examples in that the root decision is not to simply choose or filter out potential options. Instead, it involves the creation of insights assembled from evidence gathered in the tool. Interestingly, the CREATE decision while important in the decision-making problem addressed by this tool, is not directly supported by the visualization system.

VII. EVALUATION

To further validate our typology, we recruited nine visualization experts who built decision-support tools in the past to evaluate our typology and identify its strengths and weaknesses. Three of our participants are assistant professors, four are scientists at national laboratories, one is a senior visualization developer at a market research company, and one is a PhD student in visualization. The interviews were semi-structured, guided by a set of pre-determined open-ended questions (please see Fig. 7) while allowing room for deeper exploration of particular topics as they emerged. Specifically, our questions examined three dimensions proposed by Beaudouin-Lafon [73]: (1) *descriptive power*: “the ability to describe a significant range of existing [designs]”, (2) *evaluative power*: “the ability to help assess multiple design alternatives”, and (3) *generative power*: “the ability to help designers create new designs.” Each interview lasted approximately one hour and was recorded with the participants’ consent.

During each session, we introduced the typology and its properties using the car buying example, as detailed in Sections IV and V. We encouraged participants to ask questions and seek clarifications to ensure they fully understood the typology and its properties. Following this, we conducted a collaborative sketching session using shared presentation slides, where participants used the typology to describe the decision-making tasks supported by a visualization-based decision-support tool they had designed in the past. These tools varied widely in purpose, supporting tasks such as detecting fraud in financial transactions, selecting items to purchase, generating and testing data-driven hypotheses, planning urban infrastructure, and determining new bus routes, among others. As participants articulated the decision-making tasks and processes their tool supported, we collaboratively sketched corresponding diagrams (see Section V-B). Participants were then invited to review and refine these diagrams until they accurately reflected the decision flow experienced by their users. To preserve participant anonymity, we have deliberately omitted references to tools with associated publications. The remaining tools were developed in industry or are protected under non-disclosure agreements, and thus cannot be disclosed. Slides and finalized diagrams from the interviews are available in the Supplemental Material.

At the conclusion of each interview, we invited participants to provide additional feedback on the criteria used by Brehmer et al. [13] when evaluating their typology: correctness (Section VII-A), utility in the design process (Section VII-B), and comparison to other visualization typologies (Section VII-C). We

| Topic | Questions |
|---------------------------------|--|
| Intro + clarification questions | Could you provide a quick overview of your system? Please focus on what were the decision-making problems you were helping a user address. |
| | Do you have any questions about the decision-making tasks in our typology? |
| | If you had to pick one of these tasks to represent the overall decision goal(s) of your tool, which would you pick and why? (There is no right or wrong answer!) |
| Diagramming exercise | Do you have any questions about the composability of decision-making tasks? |
| | Do you think there is any decision-making composability in the problem addressed by your system? If so, can you walk us through it? If not, why not? |
| Expressiveness | How did this diagram help you describe the decision-making addressed by your tool? |
| | How representative does the diagram seem to be of the decision-making in your app? |
| Completeness and improvements | How did you feel the composition of tasks helped you express the goal of your system? |
| | Would you add any fundamental operation/capability to this diagram to more accurately describe the decision-making addressed by your tool? |
| Other typologies | How do you think this typology relates to existing typologies, taxonomies, or requirements-gathering techniques in the vis community? |
| Other | Do you have any other insights on the decision tasks and their composability you'd like to share? |

Fig. 7. Interview questions we asked our participants.

also asked them to share suggestions for improving the typology both during the interview and afterward via email if they had further insights. We discuss the limitations of the typology gathered from these interviews in more depth in Section IX.

A. Correctness

The participants discussed two topics related to the correctness of the typology: (1) how **complete** is the typology, and (2) how well the participants were able to use the typology to **express** their decision-making problems.

Completeness: In evaluating the correctness of our typology of decision-making tasks for visualization, we first assessed its completeness to determine if there were any missing decision tasks in our typology that could limit its ability to capture all aspects of decision-making. We adopted Gleicher et al.'s definition of completeness by replacing "designs" with "decision-making problems", that is, the typology needs to cover all decision-making problems we have encountered [74]. The participants in our study found that our typology was indeed comprehensive, allowing them to represent all decisions within the problem domain. As noted by one of the participants: "*I think [the typology] is a pretty good set of high level operations. [...] It took me a little bit to really clarify the distinction between activate and choose. But once I got there, it makes a lot of sense. I don't know if I can really think of another type of operation that I felt was missing. It seemed, at least in the case of my problem, a good set [of operations] to describe what's going on.*" Other participants expressed similar sentiments. Overall, all of the participants were able to describe their decision-making problems using our typology, and none of them identified any missing decision tasks or fundamental operations in our typology.

Interestingly, a few participants incorporated additional annotations when describing their decision-support tools using our typology. These annotations were used to detail the contexts of these participants' decision-making problems, including temporal information (e.g., time duration for how long a task would take or how frequently a task needs to be performed), roles (e.g., who in an organizational hierarchy should make the decision, or in a collaborative setting, which collaborator is responsible for which task), and decision styles (e.g., rational, conceptual,

analytical, etc.). Such contextual annotations complement our typology and suggest an interesting future extension of our work. We explore these concepts further in Sections IX and X.

Expressiveness: During this study, it became clear that our typology is also expressive, meaning it can be easily used to describe decision-making scenarios. All participants expressed their proficiency in composing and decomposing decisions with the typology, enabling them to construct cohesive representations of the problems their respective visualization tools addressed. They noted that the typology captured the decision-making processes inherent in problem domains and facilitated the visualization of user interactions with decision-support tools. As one of the participants said, "*[The use of the typology] really kind of helped highlight some of the complexity around the design goals that went into it. And it's definitely reflective of conversations that [my team and I] had going into the design of the tool and sort of throughout its life span. So, in that sense, I think it's very representative.*"

This sentiment is perhaps best summarized by a participant: "*Breaking down this higher level activate task...in terms of how [the user] actually interacts with the system and what the system is doing for them, I think [the typology] helped tremendously. [...] I was able to describe step by step what people were doing in the system, and those steps mapped really cleanly to the [three decision-making] tasks.*"

B. Utility in the Design Process

During the interviews, the participants discussed the potential applications of the typology. We categorize their feedback into three themes: the use of the topology for (1) **analysis and reflection** for guiding the design process of visualization tools, (2) facilitating **communication** between developers, domain experts, and end-users, and (3) **categorizing and organizing** VIS publications related to decision-making.

Analysis and Reflection of Design: Participants reported that structuring decisions using our typology diagram helped them break down decision problems, identify potential decisions that were not initially considered, and recognize what decisions lacked support in their system. In particular, the typology diagram served as an externalization of mental models for the participants to better reflect on the overall decision-making process, allowing them to step back from the development of tools and consider the bigger picture. As explained by one of the participants, "*I think this [exercise using the typology] was actually very valuable from my standpoint in reflecting on what we are building. Take a step back and reflecting on how these decisions, what we understand, and from an abstraction and typology standpoint, what they mean. So it's been actually helpful in that regard.*"

Another participant further delineated the benefits of using the typology, "*I see this as a way [...] for us to understand clearly what we are doing. Not to focus a lot on how we do it, but what we want to do. [...] Breaking down the task like that could help me to also think about new possibilities for decision-making that I didn't consider before or didn't have time to.*"

Facilitate Communication: During our interviews, participants highlighted the typology's potential to facilitate communication with domain experts and end-users during the design process by providing a higher-level perspective. One participant articulated, "If I had this taxonomy like in your work, I can better communicate with the experts. [...] I think with this taxonomy maybe we can begin from a more high level perspective, like what does the system look like or how this system would adjust its requirements from a higher perspective. And then we can drill down to each part, for creation, for activation, for choose. What kind of interface we can design for them. I think [this diagram] would facilitate this kind of brainstorming sessions."

By facilitating better communication, a participant believed that the use of the typology could improve shared understanding of the design process between the stakeholders, thereby improving the final design. "I think having some kind of decision-making type of language might help bridge that gap between 'I as the developer, designer, researcher, I think I know what you mean, I think I know what you want, and what your processes are' as opposed to the user being like 'I don't even know how to better explain this.'" Similarly, another participant compared the use of the typology to the standard practice that relies on a list of design requirements in the design process and noted, "I think this [diagram] is a better specification of the process than the list of requirements, for sure, because at least I can see here the process of the decision, not a well laid set of requirements."

Although our evaluation did not involve domain experts or end-users, our participants speculated that our typology would be intuitive for their domain experts and end-users to understand, thereby improving the design process. One participant noted, "I don't think you could really present [Brehmer et al.'s task typology] to someone outside of the [visualization] community and explain it in 10 minutes and they could like really, intuitively, get it. [Your typology] feels much more accessible." Another participant said, "I haven't practically used the diagram [with a domain expert, but] I think the expert could provide feedback directly on the diagram. Basically, if they want more alternatives, or if they think that the current way of generating alternatives is not satisfactory, they can request to add edge backward to the creation tasks. So our system design will improve." This direct engagement with the diagram has the potential to improve the collaborative refinement of system design, allowing experts to contribute insights that could enhance decision-making capabilities.

Categorization and Evaluation of Publications: Some participants highlighted the potential of our typology in aiding the research community in categorizing and organizing publications relating to decision-making. In our interviews, three participants utilized our typology to sketch diagrams representing different versions of their tools or extensions published in subsequent papers, demonstrating its capacity to compare decision processes across various tools. Through these diagrams, they were able to discuss the nuances and distinctions between these tools, highlighting areas with varying levels of decision support. This process enabled them to evaluate and compare their own tools' effectiveness.

Some participants expressed that our typology could serve as a valuable tool for evaluating decision support tools due to its ability to organize the decision process into distinct parts and delineate where decisions begin and end. One participant noted, "In my experience as a reviewer, I don't explicitly ask authors to follow specific patterns or taxonomies to organize their work. However, if they use a taxonomy, it will be easier to follow their paper. Basically, they can organize decision-making papers into these few stages, making it easier to evaluate this type of work."

Instead of delving into the minutiae of individual visualizations in the review process, one participant suggested that our typology could offer a framework to assess systems holistically, stating that, "I think [your typology] could be helpful as well when we evaluate visualizations. [In reviewing visualizations] we go very low level [evaluating if] people [can] read a certain value [from a visualization]. [...] Thinking about [a decision support] system [using your typology], like in this way, might be also helpful to evaluate the systems overall, because [as a reviewer] I know where I start and where I want to end. And I can see how [a user would] follow these steps to get to the final process."

C. Comparison With Existing Task Taxonomies

All of our participants are researchers in the VIS community who are well-versed in the VIS literature. Many drew comparisons between our typology and existing task taxonomies, such as those proposed by Amar et al. [11] and Brehmer et al. [13]. In this section, we summarize their comments on how our typology compares to existing task taxonomies in terms of: (1) the **level of task abstraction** and (2) its **flexibility** in practical use.

Level of Task Abstraction: During discussions, participants noted that our typology is more oriented towards modeling processes in contrast to the focus on data, analytical operations, and interactions emphasized by existing task taxonomies. This approach resonated with them as they considered it more suitable for describing complex decision-making tasks. One participant articulated this sentiment as follows: "I think [existing task taxonomies] can be suitable for all visual analytics systems and it's a bit low level [...]. [But] if you want to describe my system with this type of taxonomies you have to go very low level, like presenting values or like sorting the alternatives. [...] It's a bit cumbersome to describe the system using this taxonomy. Yeah, yours are better. I think your taxonomy is doing a better job at describing decision-making tasks or decision-making systems, because it's from the number of alternative perspective. So I think that the taxonomy more fits in a decision-making scenario."

Other participants also noted the challenge of using existing task taxonomies to describe a complex design-support visualization system and suggested that a higher-level typology is better suited for describing the user's needs and perspectives. As stated by one of the participants, "[Using existing taxonomies] is very low level to understand a system like [my system]. For the final users, the user is not thinking about reading this bar chart. [The] users say 'can I choose my actual scenario?'" Another participant agreed, "[Your typology] is abstract, but it helps you

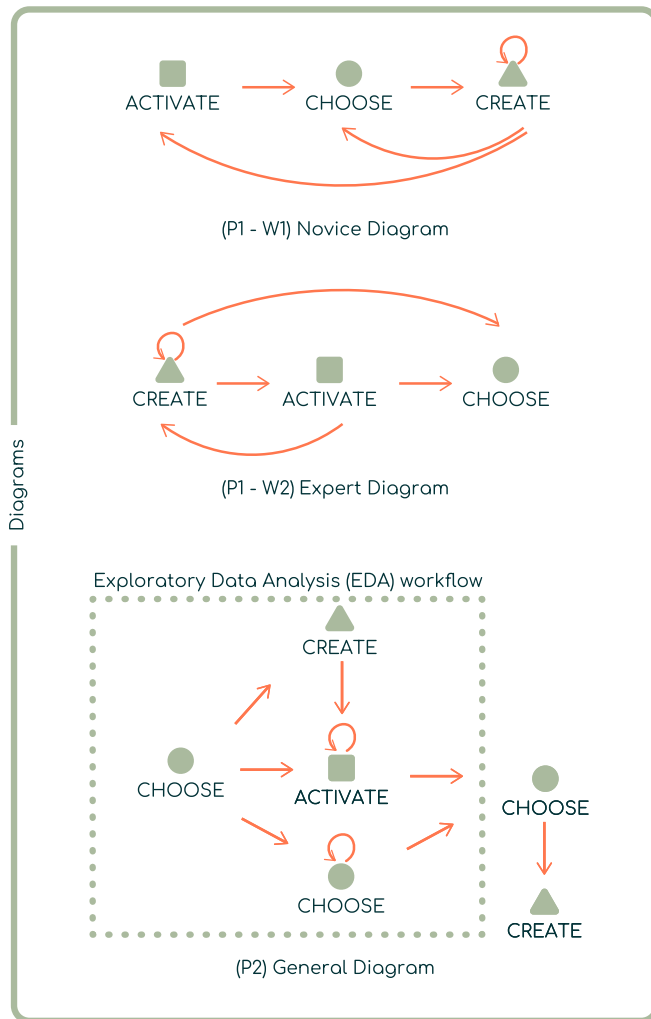


Fig. 8. Typology diagrams elicited in the first interview. The drawings have been recreated to match the paper's theme, but the content remains true to what the participant sketched during the interview.

think about the decision process. This is closer to the user than the **why** of a visualization [in the typology by Brehmer et al.], because you don't specify, for example, 'I want to present data,' or 'I want to discover data.' [...] [With a] diagram [using your typology], [...] I can see a process that I can match well with what we did in [our system]."

Flexibility to Data and Context: The participants in our evaluation highlighted the flexibility of our typology, particularly in its adaptability to different data and contexts while maintaining well-specified functions. The participants could compose and decompose decisions according to their decision tasks and preferences for modeling the problems. For example, one participant chose to represent their tool's addressed problem hierarchically, noting, "This does feel more hierarchical. I'm definitely reminded of task abstraction type of work, which tends to go a step beyond just requirements gathering."

Other participants opted for the diagram representation, utilizing cycles and arrows to depict iterations within their decision-making process. This approach allowed for clear visualization of feedback loops, with decisions feeding back into the initial

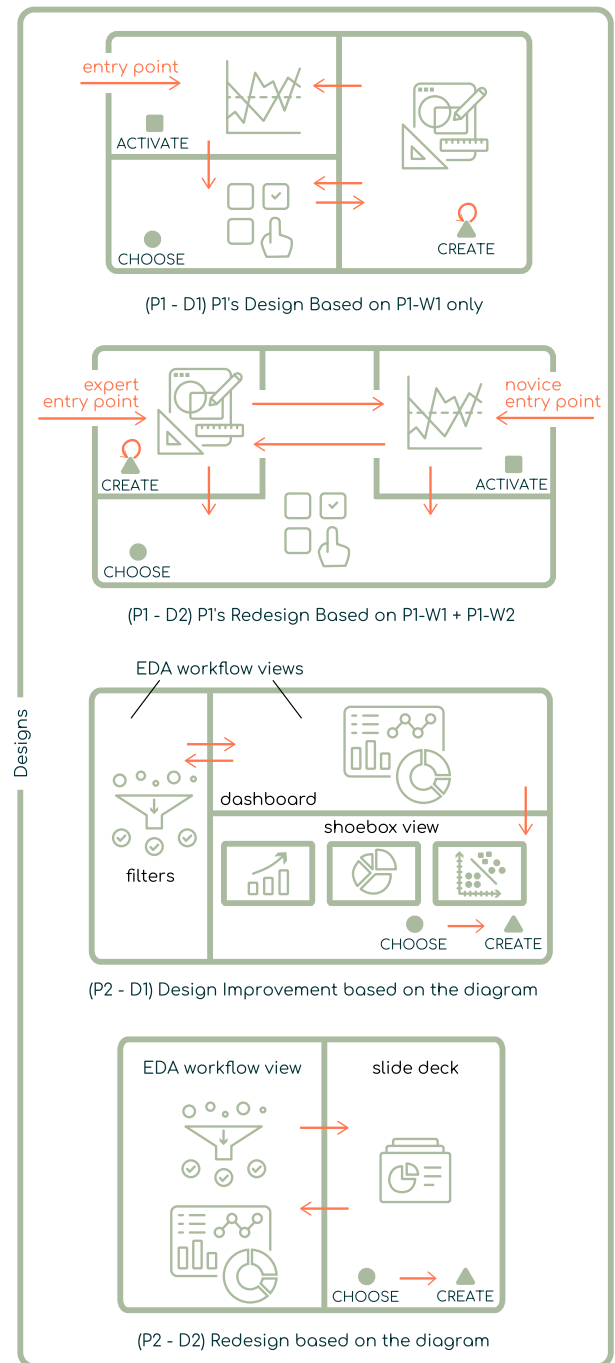


Fig. 9. Design improvements and redesigns sketched by participants during the preliminary study. Icons representing visualizations have been added to indicate coordinated views, but these were not part of the participants' original sketches.

stages of the process. One participant who adopted this method remarked, "[Brehmer et al.'s typology] very much still had kind of a rigid application to it, whereas I think the typology that you're proposing feels like it can be more flexible."

The general agreement among participants was that our typology stands out for its flexibility, allowing for diverse representations of decision-making processes tailored to specific needs and contexts. A participant summarized this quite succinctly by

stating, “*I like that a lot as part of your typology, is that it feels simple and flexible, yet powerful.*”

VIII. DESIGN IMPLICATIONS

The evaluation presented in Section VII examines the utility of our typology to describe decision-making problems supported by visualization systems. Here we present initial findings on design implications of our typology gathered from follow-up discussions with available interview participants.

Two of our nine interview participants, a PhD student specializing in design studies (P1) and a visualization developer from industry (P2), were available for a follow-up discussion about their visualization tools. We presented them with the decision-making diagrams they had created during the initial evaluation of our typology. During this session, we asked them to reflect on the diagrams and sketch potential design changes informed by the insights gained. Specifically, they were asked to: 1) identify which decision tasks were well-supported and which lacked adequate support, 2) sketch design improvements for areas where the system did not provide sufficient decision support, and 3) conceptualize a new design from scratch based on the problem outlined in the diagram. Both participants utilized our typology to enhance their designs, incorporating new views, layouts, and interactive features to better support and expand the decisions and processes within their tools.

Using the diagram created in our previous interview, P1 redesigned their tool to better accommodate both novice and expert users. The initial tool design had a workflow where numerous eligibility criteria are evaluated, filtered, and ultimately selected, leading to the creation of more restrictive criteria that feed back into the decision-making process (Fig. 8 P1 - W1). This approach catered to novice users lacking the experience necessary to create specific criteria without overlooking crucial variables. However, after creating the decision-making diagram in our first interview, P1 realized that expert users required a different workflow that allows them to create criteria from scratch, exhibited by the second diagram (Fig. 8 P1 - W2) drawn in the follow-up discussion. In the initial tool design (Fig. 9 P1 - D1), only one of these workflows is supported at a time. In redesigning the tool, the participant integrated both workflows by allowing two separate entry points into the tool (Fig. 9 P1 - D2). This redesign supports a broader range of users and allows flexibility in utilizing both workflows when needed. Furthermore, P1 used the decision-making diagrams to reflect further on the design and interaction across the views associated with each decision. This resulted in P1 adding coordinated views to facilitate the ACTIVATE and CREATE decisions, and replacing an overview visualization with more interactive capabilities to support experts in the CREATE decision. Thus, the typology diagram might inform high level design patterns, such as adding coordination between views that correspond to decisions that are linked in the diagram, or considering the information needed on the landing page of the system depending on which is the first decision that needs to be made. However, P1 hinted at the need for some correspondence from the decision-making typology and design decisions, e.g. the choice of certain visual

encodings. Exploring the mapping between decision-making tasks and design patterns is an important direction for future work.

P2 reflected on the lack of support in their visualization tool for some of the decisions contained in their decision diagram. The tool primarily supports the exploratory data analysis stage (P2’s EDA workflow in Fig. 8) where users iteratively explore data and make intertwined decisions, followed by a sequential stage that culminates in a “CREATE report” decision (the CHOOSE → CREATE motif in P2’s diagram in Fig. 8). P2 noted that the final CREATE decision is not well-supported by the tool, primarily because the initial requirements gathering focused more on the iterative, exploratory, and data-driven aspects of the process, rather than the sequential stages of decision-making. As a result, tasks like creating presentation slides are left to existing tools rather than being integrated into the system. However, P2 noted that a possible improvement in their tool that allows users to save intermediate results during exploration could better support the final CREATE decision. To address this, P2 sketched and considered two possible changes in the tool design: 1) a “shoebox” view where users drag and drop elements and insights gathered during the EDA process (Fig. 9 P2 - D1), and 2) a split view with the ability to drag and drop results directly into presentation slides (Fig. 9 P2 - D2)). P2 ultimately preferred the split view design as it provides continuous visibility and directly supports the final CREATE decision. Observing the differences between the design improvement exercise and the redesign exercise, we found the participant shifting from a bottom-up towards a top-down approach. Rather than focusing on the data or individual steps of the EDA process, more time was spent considering the overall decision goal and information needs of the user to achieve that goal.

While these are only preliminary findings, there is evidence that our typology aids in the design of visualization systems. In future work, we intend to investigate further how our typology assists in the design of decision-support visualization systems.

IX. LIMITATIONS

A few limitations of our typology emerged during our interview study: (1) the imprecise description of decision-making context, (2) the flexibility in diagram decomposition, and (3) the lack of guidance for final design decisions.

First, the typology describes decision-making tasks, but it does not account for broader contexts in which these tasks can occur. During the interview study, several participants supplemented their diagrams with annotations to provide context, without altering the core tasks or their definitions. These annotations included details such as the frequency of decision-making instances, roles within collaborative or organizational settings, and decision-making styles. While these annotations do not change the core definitions of the tasks, they often influence the design of decision-support systems. For example, one participant suggested using the typology to depict decision-making based on different “personas,” resulting in diagrams that reflect distinct roles within an organization, while another participant proposed annotating diagrams to differentiate between tasks

performed by humans and those involving computer assistance or collaboration.

Second, the typology's flexibility in decomposing decision-making tasks also presents a challenge. There are no rules or exact guidance on the appropriate level of granularity to represent decision-making problems. Our participants annotated the inputs and outputs of each decision task on the arrows that connect them. However, across participants, we have seen that the level of detail varied significantly—some provided high-level information, while others included very specific dataset details. This lack of a stopping point in the recursive decomposability of the tasks mirrors limitations seen in other typologies, such as the under specified abstraction levels found in Brehmer et al.'s typology [13]. While this flexibility can ease diagram creation, it presents challenges in maintaining consistency.

Lastly, we primarily focused on the descriptiveness of decision-making problems, and how to use these descriptions as tools in visualization design. However, our typology offers no direct mechanism to compare designs that satisfy the same decision-making diagram. We acknowledge that selecting between design alternatives is influenced by various factors beyond the typology itself, such as the designer's experience, time constraints, and the decision-making style of the domain expert, which are not captured by our typology.

While our typology offers valuable insights and flexibility in describing decision-making processes, these limitations indicate areas for refinement and further exploration. In the following section, we outline our future work, which will address some of these limitations and expand upon our typology's potential applications.

X. FUTURE WORK

While the limitations section outlined some constraints of our typology, future work will focus on leveraging the typology as a foundation for building new visualization knowledge, applications, and tools. The following points summarize our intended directions:

Full-Scale Design Evaluation: A comprehensive design evaluation with designers and their domain experts will be conducted to explore the typology's utility in end-to-end visualization design settings. This will involve longitudinal studies rather than brief interviews, allowing for deeper insights into how the typology can be applied in real-world design processes.

Expanding Design Integration: Informed by our user studies, future work will involve integrating the diagrams generated by our typology with corresponding visualization designs, potentially leading to the development of recommendation systems for decision-support tools, such as the Draco system [75]. We plan to conduct an in-depth analysis of recurring decision-making patterns, such as multi-criteria decision-making, which could evolve into an open-source project where contributors share typology diagrams. This would enable the study of decision-making patterns and the recommendation of similar problems and design solutions.

Connecting Decision-Making and Analytical Tasks: Another avenue for future work is exploring connections between decision-making tasks in our typology and established analytical tasks, such as those in Amar and Eagan's taxonomy [11] or

Brehmer et al.'s typology [13]. Identifying patterns in these connections could help predict decision-making tasks based on user interaction data or suggest relevant analytical tasks based on decision goals. This would have a direct impact on the reasoning about decision-making tasks in the context of visual analytics, as it can enhance the alignment between decision support systems and the analytical workflows that users engage with, ultimately improving the efficiency of visual analytics tools. Recent work on modeling user interactions in large-scale visualization systems [76] provides a promising foundation for this direction, as it demonstrates how structured representations of user interactions (e.g., via augmented statecharts) can be used to understand and optimize user behavior. Integrating such models with decision-making task analysis could enable systems that adapt not only to performance needs but also to users' decision intents further enhancing the effectiveness of visual analytics tools.

Human-Machine Decision-Making: We plan to use our typology to explore and reason about the division between human and machine decision-making in systems, and how designs can reflect or leverage this distinction. We believe our typology can provide insights into human-machine collaboration, the effects of automating parts of the decision-making process, and the role of human intervention in AI-driven systems. This is especially crucial in visual analytics, where effective collaboration between humans and AI ensures interpretability, trust, and control over complex decision-support tools [77].

Overall, the next steps will involve applying our typology to the development and improvement of tools and decision-making processes, seeking to validate its relevance and effectiveness in real-world design and decision-making scenarios.

XI. CONCLUSION

This paper presents a typology for decision-making tasks in visualization, addressing the limitations of existing task taxonomies. Built upon prior research and informed by design goals derived from a thorough literature review, the typology comprises three decision tasks: CHOOSE, ACTIVATE, and CREATE. These tasks allow for the representation of complex decision-making structures, as they can be composed or decomposed into other tasks. The typology demonstrates completeness, expressiveness, and utility across various decision-making contexts and is validated through semi-structured interviews with visualization experts. During the interviews, the visualization experts expressed that the typology has the potential to be used in analyzing designs of decision-support visualization systems, facilitating communication between designers and domain users, and categorizing visualization publications on decision-making. By providing a structured approach to understanding decision-making tasks, the typology offers a tool for researchers and practitioners alike, with implications for improving designs of decision-support tools in the future.

ACKNOWLEDGMENT

Funding provided by the U.S. Department of Energy Office of Science, through the Advanced Scientific Computing Research program. The views expressed in the article do not necessarily

represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

REFERENCES

- [1] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim, "Knowledge generation model for visual analytics," *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 1604–1613, Dec. 2014.
- [2] J. J. Thomas and K. A. Cook, "A visual analytics agenda," *IEEE Comput. Graph. Appl.*, vol. 26, no. 1, pp. 10–13, Jan./Feb. 2006.
- [3] C. Ware, *Information Visualization: Perception for Design*. San Mateo, CA, USA: Morgan Kaufmann, 2019.
- [4] E. Dimara and J. Stasko, "A critical reflection on visualization research: Where do decision making tasks hide?," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 1, pp. 1128–1138, Jan. 2022.
- [5] S. Afzal, R. Maciejewski, and D. S. Ebert, "Visual analytics decision support environment for epidemic modeling and response evaluation," in *Proc. 2011 IEEE Conf. Vis. Analytics Sci. Technol.*, 2011, pp. 191–200.
- [6] J. Waser et al., "Many plans: Multidimensional ensembles for visual decision support in flood management," *Comput. Graph. Forum*, vol. 33, no. 3, pp. 281–290, 2014.
- [7] R. Chang et al., "Wirevis: Visualization of categorical, time-varying data from financial transactions," in *Proc. 2007 IEEE Symp. Vis. Analytics Sci. Technol.*, 2007, pp. 155–162.
- [8] T. Alameldin and M. Shepley, "New techniques in the design of healthcare facilities," in *Proc. Conf. Vis.*, 1994, pp. 374–377.
- [9] N. Ferreira et al., "Urbane: A 3D framework to support data driven decision making in urban development," in *Proc. 2015 IEEE Conf. Vis. Analytics Sci. Technol.*, 2015, pp. 97–104.
- [10] A. Gaba et al., "My model is unfair, do people even care? Visual design affects trust and perceived bias in machine learning," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 1, pp. 327–337, Jan. 2024.
- [11] R. Amar, J. Eagan, and J. Stasko, "Low-level components of analytic activity in information visualization," in *Proc. IEEE Symp. Inf. Vis.*, 2005, pp. 111–117.
- [12] J. S. Yi, Y. ah Kang, J. Stasko, and J. A. Jacko, "Toward a deeper understanding of the role of interaction in information visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 13, no. 6, pp. 1224–1231, Nov./Dec. 2007.
- [13] M. Brehmer and T. Munzner, "A multi-level typology of abstract visualization tasks," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 12, pp. 2376–2385, Dec. 2013.
- [14] E. Dimara, H. Zhang, M. Tory, and S. Franconeri, "The unmet data visualization needs of decision makers within organizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 12, pp. 4101–4112, Dec. 2022.
- [15] E. Oral, R. Chawla, M. Wijkstra, N. Mahyar, and E. Dimara, "From information to choice: A critical inquiry into visualization tools for decision making," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 1, pp. 359–369, Jan. 2024.
- [16] L. Shi, L. Xia, Z. Liu, Y. Sun, H. Guo, and K. Mueller, "How "applied" is fifteen years of VAST conference?," in *Proc. 2023 IEEE Vis. Vis. Analytics*, 2023, pp. 121–125. [Online]. Available: <https://ieeexplore.ieee.org/document/10360912/>
- [17] B. Fischhoff and A. L. Davis, "Communicating scientific uncertainty," *Proc. Nat. Acad. Sci. USA*, vol. 111, pp. 13664–13671, 2014. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1317504111>
- [18] H. A. Simon, *The New Science of Management Decision*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1960.
- [19] J. Hullman, A. Kale, and J. Hartline, "Decision theoretic foundations for experiments evaluating human decisions," 2024. [Online]. Available: <https://arxiv.org/abs/2401.15106>
- [20] R. A. Amar and J. T. Stasko, "Knowledge precepts for design and evaluation of information visualizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 11, no. 4, pp. 432–442, Jul./Aug. 2005.
- [21] S. K. Card, J. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*. San Mateo, CA, USA: Morgan Kaufmann, 1999.
- [22] N. Andrienko and G. Andrienko, *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Berlin, Germany: Springer Science & Business Media, 2006.
- [23] L. M. Padilla, S. H. Creem-Regehr, M. Hegarty, and J. K. Stefanucci, "Decision making with visualizations: A cognitive framework across disciplines," *Cogn. Res., Princ. Implic.*, vol. 3, no. 1, Jul. 2018, Art. no. 29.
- [24] B. Fischhoff and S. B. Broomell, "Judgment and decision making," *Annu. Rev. Psychol.*, vol. 71, pp. 331–355, 2020.
- [25] K. Takemura, *Behavioral Decision Theory*. Berlin, Germany: Springer, 2021.
- [26] Q. Wang, Z. Chen, Y. Wang, and H. Qu, "A survey on ML4VIS: Applying machine learning advances to data visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 12, pp. 5134–5153, Dec. 2022.
- [27] A. Wu et al., "AI4VIS: Survey on artificial intelligence approaches for data visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 12, pp. 5049–5070, Dec. 2022.
- [28] L. Padilla, R. Fyngenson, S. C. Castro, and E. Bertini, "Multiple forecast visualizations (MFVs): Trade-offs in trust and performance in multiple COVID-19 forecast visualizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 29, no. 1, pp. 12–22, Jan. 2023.
- [29] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic, "A task-based taxonomy of cognitive biases for information visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 2, pp. 1413–1432, Feb. 2020.
- [30] E. Dimara, A. Bezerianos, and P. Dragicevic, "The attraction effect in information visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 23, no. 1, pp. 471–480, Jan. 2017.
- [31] E. Dimara, A. Bezerianos, and P. Dragicevic, "Testing the attraction effect on two information visualization datasets," Ph.D. dissertation, Inria Saclay Ile de France, 2016.
- [32] H. Hosseinpour, L. E. Matzen, K. M. Divis, S. C. Castro, and L. Padilla, "Examining limits of small multiples: Frame quantity impacts judgments with line graphs," *IEEE Trans. Vis. Comput. Graph.*, vol. 31, no. 3, pp. 1875–1887, Mar. 2025.
- [33] B. Oral, P. Dragicevic, A. Telea, and E. Dimara, "Decoupling judgment and decision making: A tale of two tails," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 10, pp. 6928–6940, Oct. 2024.
- [34] L. Padilla, H. Hosseinpour, R. Fyngenson, J. Howell, R. Chunara, and E. Bertini, "Impact of COVID-19 forecast visualizations on pandemic risk perceptions," *Sci. Rep.*, vol. 12, no. 1, 2022, Art. no. 2014.
- [35] S. Pajer, M. Streit, T. Torsney-Weir, F. Spechtenhauser, T. Möller, and H. Piringer, "WeightLifter: Visual weight space exploration for multi-criteria decision making," *IEEE Trans. Vis. Comput. Graph.*, vol. 23, no. 1, pp. 611–620, Jan. 2017.
- [36] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit, "LineUp: Visual analysis of multi-attribute rankings," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 12, pp. 2277–2286, Dec. 2013.
- [37] E. Wall, S. Das, R. Chawla, B. Kalidindi, E. T. Brown, and A. Endert, "Podium: Ranking data using mixed-initiative visual analytics," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 288–297, Jan. 2018.
- [38] T. Asahi, D. Turo, and B. Shneiderman, "Using treemaps to visualize the analytic hierarchy process," *Inf. Syst. Res.*, vol. 6, no. 4, pp. 357–375, 1995.
- [39] O. Bozorg-Haddad, B. Zolghadr-Asli, and H. A. Loáiciga, *A Handbook on Multi-Attribute Decision-Making Methods*. Hoboken, NJ, USA: John Wiley & Sons, 2021.
- [40] M. I. Henig and J. T. Buchanan, "Solving MCDM problems: Process concepts," *J. Multi-Criteria Decis. Anal.*, vol. 5, no. 1, pp. 3–21, 1996.
- [41] E. Hindalong, J. Johnson, G. Carenini, and T. Munzner, "Towards rigorously designed preference visualizations for group decision making," in *Proc. 2020 IEEE Pacific Vis. Symp.*, 2020, pp. 181–190.
- [42] A. Khannoussi, A.-L. Olteanu, C. Labreuche, and P. Meyer, "Simple ranking method using reference profiles: Incremental elicitation of the preference parameters," *4OR*, vol. 20, pp. 499–530, 2022.
- [43] A.-L. Olteanu, K. Belahcene, V. Mousseau, W. Ouerdane, A. Rolland, and J. Zheng, "Preference elicitation for a ranking method based on multiple reference profiles," *4OR*, vol. 20, pp. 63–84, 2022.
- [44] C. Kuhlman et al., "Evaluating preference collection methods for interactive ranking analytics," in *Proc. 2019 CHI Conf. Hum. Factors Comput. Syst.*, 2019, pp. 1–11.
- [45] J. Kreiser, A. Hann, E. Zizer, and T. Ropinski, "Decision graph embedding for high-resolution manometry diagnosis," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 873–882, Jan. 2018.
- [46] R. Van Pelt, R. Gasteiger, K. Lawonn, M. Meuschke, and B. Preim, "Comparative blood flow visualization for cerebral aneurysm treatment assessment," in *Proc. Int. Conf. Comput. Graph. Forum*, Wiley Online Library, 2014, pp. 131–140.
- [47] X. Wang et al., "Umbra: A visual analysis approach for defense construction against inference attacks on sensitive information," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 7, pp. 2776–2790, Jul. 2022.

- [48] P. H. Nguyen, C. Turkay, G. Andrienko, N. Andrienko, O. Thonnard, and J. Zouaoui, "Understanding user behaviour through action sequences: From the usual to the unusual," *IEEE Trans. Vis. Comput. Graph.*, vol. 25, no. 9, pp. 2838–2852, Sep. 2019.
- [49] R. Scheepens, S. Michels, H. van de Wetering, and J. J. van Wijk, "Rationale visualization for safety and security," in *Proc. Int. Conf. Comput. Graph. Forum*, Wiley Online Library, 2015, pp. 191–200.
- [50] B. A. Aseniero, T. Wun, D. Ledo, G. Ruhe, A. Tang, and S. Carpendale, "STRATOS: Using visualization to support decisions in strategic software release planning," in *Proc. 33rd Annu. ACM Conf. Hum. Factors Comput. Syst.*, New York, NY, USA, 2015, pp. 1479–1488.
- [51] X. Zhao et al., "SkyLens: Visual analysis of skyline on multi-dimensional data," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 246–255, Jan. 2018.
- [52] S. Rudolph, A. Savikhin, and D. S. Ebert, "FinVis: Applied visual analytics for personal financial planning," in *Proc. 2009 IEEE Symp. Vis. Analytics Sci. Technol.*, 2009, pp. 195–202.
- [53] R. Azuma, H. Neely, M. Daily, and M. Correa, "Visualization of conflicts and resolutions in a "free flight" scenario," in *Proc. Int. Conf. Vis.*, 1999, pp. 433–557.
- [54] J. Sorger, T. Ortner, C. Luksch, M. Schwärzler, E. Gröllner, and H. Piringer, "LiteVis: Integrated visualization for simulation-based decision support in lighting design," *IEEE Trans. Vis. Comput. Graph.*, vol. 22, no. 1, pp. 290–299, Jan. 2016.
- [55] J. Waser et al., "Many plans: Multidimensional ensembles for visual decision support in flood management," in *Proc. Conf. Comput. Graph. Forum*, Wiley Online Library, 2014, pp. 281–290.
- [56] M. Booshehrian, T. Möller, R. M. Peterman, and T. Munzner, "Vismon: Facilitating analysis of trade-offs, uncertainty, and sensitivity in fisheries management decision making," in *Proc. Conf. Comput. Graph. Forum*, Wiley Online Library, 2012, pp. 1235–1244.
- [57] T. Munzner, *Visualization Analysis and Design*. Boca Raton, FL, USA: CRC Press, 2014.
- [58] M. Le Goc, L. H. Kim, A. Parsaei, J.-D. Fekete, P. Dragicevic, and S. Follmer, "Zoooids: Building blocks for swarm user interfaces," in *Proc. 29th Annu. Symp. User Interface Softw. Technol.*, 2016, pp. 97–109.
- [59] G. Campitelli and F. Gobet, "Herbert simon's decision-making approach: Investigation of cognitive processes in experts," *Rev. Gen. Psychol.*, vol. 14, no. 4, pp. 354–364, 2010.
- [60] Y. Wang and G. Ruhe, "The cognitive process of decision making," *Int. J. Cogn. Inform. Natural Intell.*, vol. 1, no. 2, pp. 73–85, 2007.
- [61] E. Dimara, A. Bezerianos, and P. Dragicevic, "Conceptual and methodological issues in evaluating multidimensional visualizations for decision support," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 749–759, Jan. 2018.
- [62] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, pp. 1–35, 2021.
- [63] D. Danks and A. J. London, "Algorithmic bias in autonomous systems," in *Proc. Int. Joint Conf. Artif. Intell.*, 2017, pp. 4691–4697.
- [64] L. Battle and A. Ottley, "What do we mean when we say "insight"? A formal synthesis of existing theory," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 9, pp. 6075–6088, Sep. 2024.
- [65] R. Chang, C. Ziemkiewicz, T. M. Green, and W. Ribarsky, "Defining insight for visual analytics," *IEEE Comput. Graph. Appl.*, vol. 29, no. 2, pp. 14–17, Mar./Apr. 2009.
- [66] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon, *Visual Analytics: Definition, Process, and Challenges*. Berlin, Germany: Springer, 2008.
- [67] N. Chinchor and W. A. Pike, "The science of analytic reporting," *Inf. Vis.*, vol. 8, no. 4, pp. 286–293, 2009.
- [68] D. J. Kasik, D. Ebert, G. Lebanon, H. Park, and W. M. Pottenger, "Data transformations and representations for computation and visualization," *Inf. Vis.*, vol. 8, no. 4, pp. 275–285, 2009.
- [69] H. A. Simon, "Background of decision making," *Nav. War College Rev.*, vol. 10, no. 3, pp. 1–24, 1957.
- [70] D. Weng, H. Zhu, J. Bao, Y. Zheng, and Y. Wu, "Homefinder revisited: Finding ideal homes with reachability-centric multi-criteria decision making," in *Proc. 2018 CHI Conf. Hum. Factors Comput. Syst.*, 2018, pp. 1–12.
- [71] M. V. Dowling, "Semantic interaction for symmetrical analysis and automated foraging of documents and terms," Ph.D. dissertation, Virginia Tech, 2020.
- [72] R. Chang et al., "Scalable and interactive visual analysis of financial wire transactions for fraud detection," *Inf. Vis.*, vol. 7, no. 1, pp. 63–76, 2008.
- [73] M. Beaudouin-Lafon, "Designing interaction, not interfaces," in *Proc. Work. Conf. Adv. Vis. Interfaces*, 2004, pp. 15–22.
- [74] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts, "Visual comparison for information visualization," *Inf. Vis.*, vol. 10, no. 4, pp. 289–309, 2011.
- [75] D. Moritz et al., "Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco," *IEEE Trans. Vis. Comput. Graph.*, vol. 25, no. 1, pp. 438–448, Jan. 2019.
- [76] D. Benvenuti, M. Filosa, T. Catarci, and M. Angelini, "Modeling and assessing user interaction in big data visualization systems," in *Proc. IFIP Conf. Hum.-Comput. Interaction*, Springer, 2023, pp. 86–109.
- [77] S. Monadjemi, M. Guo, D. Gotz, R. Garnett, and A. Ottley, "Human-computer collaboration for visual analytics: An agent-based framework," in *Proc. Conf. Comput. Graph. Forum*, Wiley Online Library, 2023, pp. 199–210.



Camelia D. Brumar received the MS degree from Tufts University. She is currently working toward the PhD degree in computer science with Tufts University. Her research focuses on designing abstractions that help structure and reason about decision processes in visualization. Her research interests include problem definition in visualization workflows, frameworks for decision-making, and how these abstractions shape the design of visualization decision-support systems.



Sam Molnar received the PhD degree in computer science from the University of Colorado Boulder. They are a research scientist with the National Renewable Energy Laboratory. Their research work focuses on developing novel visualization designs and interaction techniques to support power system stakeholders' analyses for many scenarios such as real-time operations, model comparisons, state estimation, science communication, etc.



Gabriel Appleby received the PhD degree in computer science with Tufts University. He is a research scientist with the National Renewable Energy Laboratory. His research spans the fields of data visualization, visual analytics, and machine learning.



Kristi Potter (Member, IEEE) received the PhD degree in computer science from the University of Utah. She is a group manager and scientist with the National Renewable Energy Laboratory. Her research is focused on methods to improve visualization techniques by adding qualitative information regarding reliability to the data display. This work includes researching statistical measures of uncertainty, error, and confidence levels, and translating the semantic meaning of these measures into visual metaphors.



Remco Chang (Member, IEEE) received the PhD degree in computer science from the University of North Carolina Charlotte. He is an associate professor in computer science with Tufts University. His research interests include visual analytics, information visualization, human computer interaction, and databases.