

Task Cube: A Three-Dimensional Conceptual Space of User Tasks in Visualization Design and Evaluation

Alexander Rind,
Wolfgang Aigner,
Markus Wagner
St. Poelten University of
Applied Sciences, Austria

Silvia Miksch
Vienna University of
Technology, Austria

Tim Lammarsch
MODUL University Vienna,
Private University,
Austria

ABSTRACT

User tasks play a pivotal role in visualization design and evaluation. However, the term ‘task’ is used ambiguously within the visualization community. In this article, we critically analyze the relevant literature and systematically compare definitions for ‘task’ and the usage of related terminology. In doing so, we identify a three-dimensional conceptual space of user tasks in visualization, referred to as the task cube, and the more precise concepts ‘objective’ and ‘action’ for tasks. We illustrate the usage of the task cube’s dimensions in an objective-driven visualization process, in different scenarios of visualization design and evaluation, and for comparing categorizations of abstract tasks. Thus, visualization researchers can better formulate their contributions which helps advance visualization as a whole.

Keywords

Action, design guidelines, interaction, objective, task frameworks, terminology, visualization theory

1. INTRODUCTION

Tasks are an issue discussed frequently in the visualization community as they are pivotal to how we design and evaluate our work. In many, if not all, scenarios of empirical visualization research [25; 19] tasks are either central to the study setup (e.g., in controlled experiments or case studies) or an emergent study result (e.g., transcription of observed visual data analysis and reasoning processes). This is further illustrated by the role of tasks in widely used design frameworks such as the Data–Users–Tasks Design Triangle [34] or the Nested Model [35].

However, there is continuing confusion about what the term ‘task’ means in a visualization context. Even if we consider tasks only for visualization users and neither for developers of visualization systems nor for the audience of a presentation, there are many nuances of what such a visualization task can be. It may be as open-ended as ‘detect

anomalies in recent public health data’ or ‘identify the main drivers of climate change’ but also as crisp as ‘find yesterday’s most profitable product’ or “buy a train ticket” [56, p. 2433]. Already in 1994 it was widely acknowledged that “the notion of ‘task’ is increasingly difficult to pin down” [44, p. 410] and nowadays the word ‘task’ is still regarded as “deeply overloaded in the visualization literature” [35, p. 921]. Based on our own experience throughout multiple visualization design and evaluation projects and inputs from fellow researchers, we regard this confusion as unsatisfactory. A commonly agreed understanding and terminology of tasks are needed.

Therefore, this article investigates user tasks in visualization design and evaluation using the three conceptual dimensions, referred to as the task cube, as a theoretical compass:

- The primary contribution of this article is the *conceptual space* and its dimensions perspective, abstraction, and composition, which we define and describe in the section of the same title. We also compare our dimensions systematically to definitions from relevant literature in visualization and human-computer interaction (HCI).
- Furthermore, we emphasize the central *role of tasks in visualization* in the next section. On the one hand, we present visualization as a task-driven endeavor by putting objectives and actions to the foreground of visualization processes. On the other hand, we illustrate the role of tasks in different scenarios of visualization design and evaluation.
- Finally we survey 37 *categorizations of abstract tasks* found in theoretical frameworks and state-of-the-art reports and systematically compare these categorizations along the dimensions of the task cube.

Thus, instead of proposing a categorization of individual tasks, the task cube and its dimensions help visualization researchers navigate the conceptual space of task categorizations – in particular when working on theoretical task frameworks and collecting state-of-the-art reports. In addition, the reduced ambiguity allows visualization researchers to better formulate their contributions. This manuscript builds on contributions presented in a previous workshop article [47] with revised definitions, an objective-driven visualization process, an illustrative example, and a survey of abstract task frameworks as major extensions.

The final, definitive version of this paper has been published in **Information Visualization** by SAGE Publications Ltd, All rights reserved. © The Authors

Accepted for publication: November 5, 2015
Published online before print: December 27, 2015
doi: [10.1177/1473871615621602](https://doi.org/10.1177/1473871615621602)




Figure 1: Overview of the conceptual space of user tasks in visualization with composition, perspective, and abstraction as orthogonal dimensions.

2. CONCEPTUAL SPACE

In an effort to clear the aforementioned confusion around the term ‘task’, we analyzed the scientific literature in HCI and visualization for definitions of what tasks are and what roles they play in visualization (see section on *Comparison to Literature*). In addition, we surveyed and compared task categorizations published in theoretical frameworks and state-of-the-art reports (see section on *Categorizations of Abstract Tasks*). Our materials have a focus on visualization and visual analytics of time-oriented data and graph data. Visual analytics [63] has a particular emphasis on developing a science of interaction [2] and to support higher-level reasoning. Nevertheless, we consider our results applicable to the visualization field in general.

Our literature review confirmed that the term ‘task’ is not consistently defined but refers to different concepts. Consequently, our analysis concentrated on characterizing these concepts and we identified three dimensions – abstraction, composition, and perspective – spanning the task cube (Figure 1). For this purpose, we treat composition and abstraction as continuous dimensions along with the why/how dichotomy for perspective. Thus, a concept of task can be placed either on the top or on the bottom plane of the figure cube. As we demonstrate below these dimensions are nonredundant and sufficiently expressive to explain the conceptual variety of user tasks in different visualization design and evaluation scenarios.

Next we present these dimensions in a consolidated terminology distilled from our literature analysis:

Abstraction. On the one hand, *concrete tasks* describe what needs to be done in a specific application context, such as ‘find the quarter of Google’s largest revenue’. On the other hand, *abstract tasks* describe an aim on a more generic level, such as ‘find maximum’. To generalize the concept behind a concrete task, they can be abstracted and expressed using generic categories from task frameworks. For example, Andrienko and Andrienko [7] would classify this as the abstract task ‘indirect lookup’. Thus, abstraction and task frameworks allow systematic study and facilitate the reuse of visualization methods for tasks of the same abstract category.

Composition. Tackling a task, it is common practice to break it down into smaller subtasks [16; 44]. Thus, the level of composition can range from long-term challenges like ‘end poverty’ to small steps like ‘find outliers in economic data’. By the same token, task frameworks often distinguish

high-level, low-level, and sometimes levels in-between [5; 9; 48]. The *compositions* of tasks form hierarchies for which we propose a continuous scale from *high-level* to *low-level*.

Abstraction can be clearly distinguished from composition: a low-level task is a part of a high-level task, whereas an abstract task is a more generic category of a concrete task. This is illustrated as a low-level task demands fewer steps to perform than a high-level task encompassing it. In contrast, if an abstract task describes a given concrete task more generically, both tasks still require the same steps. These dimensions conform to the subclass-of and part-of relationships found e.g. in object-oriented programming and ontology languages.

Perspective. On the one hand, high-level visualization tasks are typically non-routine objectives such as ‘find suspicious patterns’ or ‘what explains this behavior’. For these, the technical term ‘objective’ is adopted from Roth [48] because other terms such as ‘problem’ or ‘goal’ are understood differently in our community. We define it as:

An *objective* is a question on data that the user raises for inquiry, consideration, or solution while aiming to solve a problem or satisfy a curiosity.

On the other hand, tasks using interactive features of visualization artifacts and categories of some low-level task frameworks are formulated as executable actions such as ‘map time to y-axis’ or ‘zoom to the orange cluster’. The technical term ‘action’ is widely used and we derive its definition from Gotz and Zhou [17]:

An *action* is a discrete step towards addressing an objective.

Since there is usually no direct relation – no decomposition – between objectives and actions, it makes sense to distinguish between a *why* perspective and a *how* perspective and to have two distinct terms to explicitly address these perspectives.

Both objectives and actions can be composed or decomposed at different levels: On the *why* perspective, users break down large objectives to increasingly smaller *subobjectives* intentionally or unintentionally in order to make ill-defined objectives manageable using visualization tools [16; 38]. On the *how* perspective, blocks of consecutively observed actions can be combined to *action sequences*, at multiple levels of composition. With experience, users will develop *action strategies* to solve common subobjectives with a tool and it is possible to analyze such strategies by observing users [43]. In addition, actions can be decomposed further to a level below intentional problem solving such as individual user interface events [17]. Even though the composition hierarchies of objectives and actions are connected, it makes sense to draw a clear line between them along the *why/how* dichotomy because objectives and actions are different concepts with different properties. Furthermore, actions can be found at the same composition level as intermediate- to low-level objectives with the mapping varying largely depending on visualization tools. For example a specialized tool for electronic health records might have the affordance to select all critical patients by a single intermediate-level action whereas a general purpose time series analysis tool requires a sequence of low-level actions to achieve the same objective.




Figure 2: Visual task decomposition of the Kronos Incident as described in the *Example* section.

Neither could we subsume perspective under abstraction because both objectives and actions can be described at different levels of abstraction. Above, we mentioned the concrete objective ‘find the quarter of Google’s maximum revenue’ and classified it as the abstract objective ‘indirect lookup’. Likewise, the action of changing the y-axis of a scatter plot to a different variable can be abstracted as ‘re-configure’ [70], ‘arrange’ [9], or ‘visualize’ [18].

Together these three dimensions span the task cube (Figure 1), a simple yet expressive model of the conceptual space for user tasks in visualization.

2.1 Example

To illustrate these concepts more deeply, we will follow an extensive example based on the Kronos Incident, the fictional application context created for VAST Challenge 2014 [12].

The *concrete, high-level* objective of the Kronos Incident [12] is to support law enforcement agencies of the fictional country of Kronos in investigating the disappearance of some employees of GASTech, a natural gas production company. The VAST Challenge breaks the objective down to subobjectives of (1) reconstructing the events of the disappearances, (2) reconstructing networks of people that influenced each other, (3) finding missing information, and (4) providing the best way of action for the police force (Figure 2). Of course, these *concrete* subobjectives need to be broken down further. For example, to specifically analyze the network of the main suspect, an organization named Protectors of Kronos (POK), we can identify *lower-level* objectives as (2a) identifying the leaders, (2b) identifying all members of the extended network, (2c) analyzing the change of the network structure over time, and (2d) finding potential connections between the POK and GASTech.

To break down the *concrete* subobjective (2a) (identifying the leaders of POK) even deeper, we follow the solution of Saraf et al. [50], winners of the Grand Challenge Award. They first (2a1) located co-occurrences of the terms “POK” and “leader” in news articles over time. This resulted in three time ranges when both terms peaked simultaneously. Then they (2a2) browsed through word clouds of persons retrieved from named entity recognition of articles in these time frames.

It is possible to *abstract* the high-level objective by domain as a law enforcement objective. The subobjective (2a) identifying the leaders can be *abstracted* as ‘discover/explore/

compare’ using the three-level *why*-typology by Brehmer and Munzner [9]. We can abstract the *low-level* objective (2a1) of finding co-occurrences of the terms over time as ‘discover/locate/summarize’ [9] or as a ‘synoptic inverse comparison task’ [7]. The following *concrete* objective of finding the most mentioned name in each time frame can *abstractly* be considered an ‘elementary comparison’ task according to Andrienko and Andrienko [7].

Saraf et al. [50] solve the *concrete low-level* objective (2a1) by entering the terms in a search box of their visualization artifact and then use the date range widget for subobjective (2a2). Both *concrete* actions can be *abstracted* as ‘filter’ in the frameworks of Yi et al. [70] and Gotz and Zhou [17].

2.2 Comparison to literature

In colloquial English a ‘task’ is understood as “a piece of work that has been given to someone : a job for someone to do”.¹ The Merriam-Webster Dictionary emphasizes that tasks are characterized as being externally assigned, having a deadline, and being hard or unpleasant. Surveying the HCI and visualization literature we can however discover a shifted focus and a multitude of nuances of this general definition. Next, we analyze these along the three dimensions presented above.

Abstraction. The need for abstraction and categorization of tasks is widely acknowledged. There are many different terms used for abstract tasks in the literature including ‘generic’, ‘general’, but also ‘high-level’. Munzner [35, p. 921] distinguishes between a task formulated in terms of an application domain, which she calls ‘problem’, and an abstract task, which she denotes as ‘operation’. Apart from that we are not aware of any term for a concrete task.

Composition. It appears as a general notion that users break down tasks into smaller, better manageable *subtasks* [16]. Preece et al. [44, p. 412] speak of a “decomposition of goals into subgoals and tasks into subtasks as the user moves downwards through a hierarchy of systems”. One possible notation to specify such hierarchical tasks models are ConcurTaskTrees [41]. In the visualization literature, the terms ‘high-level’ and ‘low-level’ are often used to distinguish such tasks [9] but there is no consistent understanding.

In addition, some task frameworks introduce explicitly named composition levels (Figure 3), which contrasts abstraction that to our knowledge has no such hierarchies. Norman [39, p. 15] introduces a composition level above tasks: “an activity is a coordinated, integrated set of tasks” and emphasizes that design should be centered on these activities. Gotz and Zhou [17, p. 46] describe how “analysts typically follow a divide-and-conquer approach, performing several sub-tasks to achieve the requirements of a single top-level task”.

Tasks at the leaf-level of the hierarchy are often treated distinctively. Preece et al. [44, p. 411] define an *action* “as a task that involves no problem solving or control structure component” and is performed by “the human physically interact[ing] with a device”. Tasks at a similar level of composition are also referred to as ‘simple tasks’ or ‘unit tasks’ [44]. Fuchs [16, p. 10] uses ‘basic tasks’ for the leaf-level and differentiates them from ‘actions’, which “describ[e] the functional properties beyond the conceptual task decompo-

¹<http://www.merriam-webster.com/dictionary/task>, accessed Aug 27, 2014.




Figure 3: Comparison of composition levels in different hierarchies. The grey area denotes whether the level is regarded specific to the used tool.

sition”. He defines an ‘action’ as “an atomic operation that is executed upon an artifact, by an entity that is involved in the completion of the task (user, computer, . . .)”. Similarly, the action tier of Gotz and Zhou [17, p. 46] represents “an atomic analytic step performed by a user with a visual analytic system”. They consider yet another level below actions: “events correspond to the lowest-level of user interaction events (for example, a mouse click or a menu item selection) which carry very little semantic meaning” [p. 43]. Norman [39, p. 15] describes tasks as “composed of actions, and actions [as] made up of operations”. ConcurTaskTrees [41] distinguish at the leaf-level between ‘user tasks’, ‘application tasks’, and ‘interaction tasks’. They refer to tasks composed of such subtasks as ‘abstract tasks’.

The selection of tools also plays an important role in task decomposition (Figure 3). Both Preece et al. [44] and Cooper et al. [13] use it to distinguish between ‘goals’ and ‘tasks’. For Fuchs [16] it is the ‘action’ that involves an artifact such as a task-specific visual representation.

Perspective. The distinction between a *why* and a *how* perspective of tasks stretches across the literature and has only recently been addressed by Brehmer and Munzner [9] as ends-means ambiguity.

On the one hand, the *why* perspective describes a task by the ends or the intended outcomes. Preece et al. [44, p. 411] define an external task or goal “as a state of a system that the human wishes to achieve”. Roth [48, p. 2357] distinguishes further between “an ill-defined task, or goal, motivating use of the visualization” and “a well-defined task, or objective, that can support the goal”. In visualization, the *why* perspective is often formulated as a question or a query to be answered based on the data at hand. Andrienko and Andrienko [7, p. 49] “use the word ‘tasks’ to denote typical questions that need to be answered by means of data analysis”, Amar et al. [5] categorize tasks by questions or queries asked, and Munzner [35, p. 921] denotes a task described in domain terms as ‘problem’. In her book [36, p. 45], she refers to abstract *why* tasks as ‘action’. For Gotz and Zhou [17, p. 46] the “task tier captures a user’s high-level analytic goals”.

On the other hand, the *how* perspective addresses the means or actions by which users perform the task. Preece et al. [44, p. 411] define a task, in particular an internal task, “as the activities required, used or believed to be necessary to achieve a goal using a particular device”. Fuchs [16, p. 10] understands a task “as a single, conceptually distinguishable but not necessarily atomic step within a composite activity

or work flow”. Cooper et al. [13, p. 15] write “both activities and tasks are intermediate steps (at different levels of organization) that help someone to reach a goal or set of goals”. For Schulz et al. [52, p. 2366] visualization tasks are “activities to be carried out interactively on a visual data representation for a particular reason”, for Brehmer and Munzner [9, p. 2376] “abstract tasks are domain- and interface-agnostic operations performed by users”, and Roth [48, p. 2357] distinguishes a third category: “a system function, or operator, that may support the objective”.

Discussion. What makes the notion of ‘task’ so hard to understand are not only the differences in these three dimensions (Figure 1) but also their interconnectedness. For example, when we follow a concrete user session, we could observe how the user transforms a high-level *why* task formulated in terms of the application domain into low-level *how* tasks matching operations of the visualization system [17]. Thus, it is difficult to tackle tasks in a more abstract way needed to draw generalizable guidelines and to characterize tasks on the levels in-between.

Bridging between the extremes and taking an abstract view on an intermediate level appears as a promising direction. Yi et al. [70] propose to categorize actions by user intents, i.e., “what users achieve by using the interaction techniques rather than how the techniques provided by Infovis systems work”. Schulz et al. [52] establish a five-dimensional design space of visualization tasks that encompasses both *why* and *how* as well as dimensions pertaining to the data (characteristics, target, cardinality). Likewise, Brehmer and Munzner [9] represent a task as a triple composed of *why*, *how* and *what* (data).

In contrast, our work does not aim for such a single intermediate level that captures all diverse forms of tasks. Instead, it emphasizes the diversity of task concepts. The task cube can accommodate for different task categorizations that suit different visualization scenarios and application domains. Thus, it proceeds from previous work considering its dimensions: On the one hand, Munzner [35] proposes a 2x2 matrix of task concepts by abstractions and composition. On the other hand, Roth [48] distinguishes between different concepts for *why* (goal, objective) and *how* (operator). Pike et al. [42] combine the *why/how* dichotomy and the level of composition. They describe “analytic discourse as the relationship between [interaction] techniques and the user’s goals and tasks, which involve low-level choices about manipulating interactive controls and higher level goals surrounding the problem being investigated” [42, p. 265]. Brehmer and Munzner [9] propose three dichotomies to compare task categorizations: level (high vs. low), temporality (sequences vs. constraints), and specificity (general vs. specific). What the task cube adds is a more fine-grained understanding of abstraction and a terminology that can distinguish between *why* and *how* tasks. In particular, we need to emphasize that visualization often addresses open-ended objectives. When solving such objectives, the users can follow different strategies and there are often no definite mappings between the *why* and the *how*. Furthermore, HCI scholars have long warned against a design approach that focuses on hierarchical analysis of operational tasks and not on the goals and characteristics of users [8; 13; 39; 44].

Next, we will explore the role of objectives and actions both in a concrete process and in different scenarios of design and evaluation.

3. ROLE OF TASKS IN VISUALIZATION

We conceptualize visualization-supported data analysis as a process that is largely driven by objectives. This process combines and extends existing process models by expanding the gulf of goal formation [24; 38] beyond actions towards higher-level objectives and by adding emphasis on objectives and their breakdown to the knowledge generation model [49] and the hypothesis-driven model by Lammarsch et al. [26]. For this purpose we rely on a few technical terms *model*, *finding*, *insight*, *hypothesis*, and *knowledge* from the knowledge generation model [49].

This visualization-supported data analysis process (Figure 4) starts with a user, data, and an objective [34]. While the data resides in a more or less structured form within a computer, the user brings in the objective and his or her background knowledge [49; 56]. The objective can originate from the user’s personal goals, from the goals of the user’s organization [44], or out of pure curiosity [61].

Under consideration of her or his *background knowledge* about the objective, about the data, and about the available tools, the user breaks the objective down to manageable *subobjectives*. Then he or she develops a *plan*, which is a sequence of *actions* that she or he believes will solve the subobjective and thus will advance towards addressing the overall objective [44; 16]. This planning is guided but also limited by the affordances of the tool(s) used. Such a tool is typically a visualization artifact providing interaction with visual representations and models of data. Nevertheless, we should keep in mind that in many realistic data analysis application contexts users have a range of possible tools at their disposal: e.g., other computing systems, pen and paper, or mental models.

The user performs the planned actions by activating a number of *events* in the visualization user interface that are interpreted by interaction techniques [17]. For example, panning using a scrollbar involves a mouse-down, several mouse-drag, and a mouse-up event. Still, we focus on actions as the discrete steps that represent the lowest level of activity of which the user is consciously aware. Above actions we consider *action sequences* and *action patterns* as frequently occurring sequences, e.g., sort, then select top 10. With experience, users can learn strategies involving such action patterns as generic and low-level plans, to solve common subobjectives with a tool or family of tools [43].

While performing these actions, the user collects *findings*, potentially interesting observations on the data. *Insights* arise from accumulating these findings and setting them in context with the objective and knowledge [49]. These insights often lead to an update of the plan as the user develops hypotheses that make new subobjectives relevant. Furthermore, the plan will be iteratively updated as the background knowledge, the availability of tools, or even their goals can change [44; 37; 38; 13; 42].

3.1 Implications for design and evaluation

Above, we observed, from the users’ point of view, concrete objectives and actions in the visualization process. While users can abstract their objective breakdown and action planning from one situation to another and from one visualization artifact to another, abstraction plays an even more important role for visualization design and evaluation. Here, theoretical frameworks serve as bond between individual research contributions towards general guidelines for the

visualization discipline. Next, we will walk through different scenarios of visualization design and evaluation, describe the role of objectives and actions, and give examples.

Domain characterization and abstraction. An essential aspect for understanding of environments and work practices [25] is to figure out the *objectives* of domain experts. As *high-level objectives* might be too vague, it is often necessary to gather objectives at multiple levels of *composition*. Subsequently, it is necessary to transform *concrete objectives* formulated in domain language to *abstract objectives* that can be matched to categories of task/objective frameworks. The two outer layers of Munzner’s Nested Model [35] describe these steps.

RelEx, an example study conducted in automotive engineering [55] describes objectives at three levels of composition: the high-level objective of “traffic optimization on the signalpath network” in domain terms, mid-level objectives being abstracted to general features of social networks, and low-level objectives as queries on network relations. Another study in IT security [67] classifies the objective of malware pattern analysis along the three-level *why*-typology by Brehmer and Munzner [9]. Both studies analyze objectives in context of the data and users as proposed by Miksch and Aigner [34].

Designing visualization artifacts. The next layer of visualization design [35] is mapping the abstracted *objectives* to visual representation and interaction techniques, and thus the *action* affordances provided by the visualization artifact. This step can profit from guidelines translating between abstract objectives and abstract actions [33]. An example of such design guidance for visualization of spatial and temporal data is the book of Andrienko and Andrienko [7], which lays on the solid foundation of their theoretical framework of low-level objectives.

The RelEx design study [55] illustrates how design rationale is based on clearly formulated user objectives. Based on the objectives cited above, they analyzed existing tools and postulated design requirements – in particular the requirement for rich exploration of complex, multi-way relations. RelEx meets this requirement with brushing and linking in multiple coordinated views, especially a signal-path view for 4-way relations. Likewise, the LiveGantt design study [22] tackles the objective of what-if questions in scheduling by drag and drop actions in its Gantt chart view. The ChronoLenses technique [71] addresses two objectives (T1) single data stream transformation such as remove trend and (T2) cross-data stream analysis such as compare two streams. These two tasks result in two types of interactive lenses.

Stimuli for experiments. User experience or performance can be evaluated by letting test users interact with visualization artifacts in a controlled environment. The stimuli presented to the test users are typically *objectives* and can be formulated as question like ‘which <items> fulfill <...>’, as imperative like ‘identify <items> that fulfill <...>’, or even prescribe the answering method like ‘click on <items> that fulfill <...>’. As argued above, the evaluation of pre-defined action sequences like ‘filter by <...>, then zoom to <...>’ is less relevant for user studies in visualization.

Depending on the study’s hypotheses, a suitable composition level must be found because *low-level objectives* can be too trivial and *high-level objectives* too open-ended for

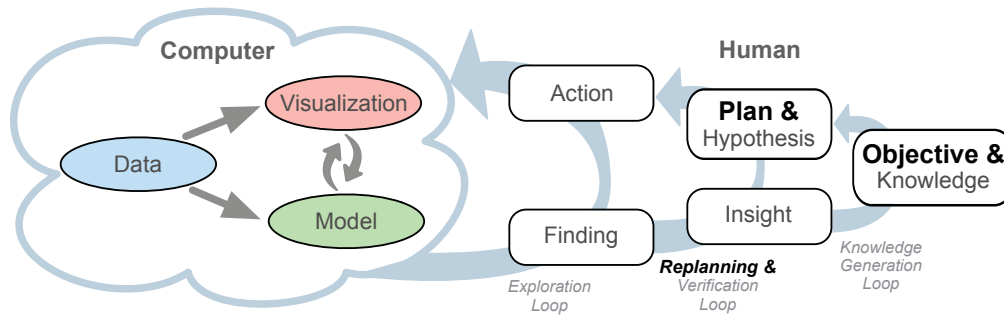


Figure 4: Objectives and actions in visualization-supported data analysis. The knowledge generation model [49] is extended with the notion of an *objective* residing at the human side, which is tackled through a *plan* and performed through a sequence of *actions*. New insights might lead to *replanning*.

quantitative analysis of time and errors. For such experiments other evaluation methods, e.g., qualitative analysis of insights [51; 60], are more suitable.

Research hypotheses are often specific to an *abstract objective* such as ‘less errors for comparison’ whereas stimuli need to present *concrete objectives* that act as representative examples for the abstract objective. In some cases a concrete objective is also translated to a concrete objective in a different domain so that a sufficiently large population of test persons can be recruited. Each occasion a test person works on such a concrete objective is called *trial* [15, p. 298]. Repeated trials with multiple test persons or multiple data (sub)sets are needed to reach statistical power. For example, Javed et al. [21] experimentally compared four visualization techniques for time series by the performance of solving comparison, slope, and discrimination objectives. The SemanticTimeZoom technique [3] was evaluated using 12 objectives, which were categorized based on the Andrienko and Andrienko [7] objective framework.

Setting context in case studies. When case studies are performed to evaluate visualization artifacts under more realistic conditions it is necessary to describe the *objective* tackled and it is useful to set it in context using *abstract objectives*. This also applies to evaluation by algorithmic performance measures or qualitative result inspection [19].

Studying interactive analysis and reasoning. Interaction remains a topic of great interest in visualization currently. Such studies include case studies with domain experts where observation of user interaction allows to see behind the curtain of visual data analysis and reasoning processes. Interaction logs can be a valuable addition to thinking aloud protocols or diaries or even replace them [25]. From sequences of *actions* in these logs, interaction patterns can be identified and generalized based on *abstract actions*. Thus, it is possible to develop shortcuts for frequently used patterns [42], explore how visualization artifacts induce or inhibit certain analytical behaviors, and judge the cost of interaction [24]

For example, the experiment of Dou et al. [14] demonstrated to what extent analytical reasoning can be inferred from interaction logs alone. They noted that such inference worked best for highly interactive artifacts but worse when users got along with looking at visual representations. Another example study by Pohl et al. [43] compared the interaction logs from user studies of two visualization artifacts of the same target audience, medical doctors, to identify in-

teraction patterns up to a length of three actions and transition probabilities between actions. In order to make actions comparable, they used the Yi et al.’s user intents [70] for abstraction.

Integration in visualization artifacts. Objectives and actions are not only relevant for visualization designers and researchers, but they can also be made explicit within visualization artifacts for users during visualization runtime thus addressing the visual analytics challenge of “capturing user intentionality” [42]. The record of *actions* performed makes sense as history mechanism [57] and for analytical provenance [40]. Notable examples are the p-set framework [20], which focuses on parameter sets, and HARVEST [17], which builds on an action taxonomy for interoperability. Likewise, the visualization artifacts can support users by allowing them to track open subobjectives as demonstrated in Sandbox [69] or Aruvi [59]. Given the user’s *objectives* the visualization artifact can also recommend suitable settings at runtime. For example, Tominski et al. [64] propose to choose one of eight color schemes based on a three-dimensional objective space and Schulz et al. [52] use their five-dimensional task framework to recommend visual representation techniques for climate impact data.

4. CATEGORIZATIONS OF ABSTRACT TASKS

Beyond individual visualization design and evaluation projects, tasks are useful for systematic research in visualization. Thus, a number of theoretical frameworks that categorize abstractions of objectives and actions have been conceived. These frameworks are valuable to make the results of empirical research comparable and to extract design guidelines from visualization work.

State-of-the-art reports, in particular, can apply such categorizations to describe and systematically compare visualization artifacts. The objectives addressed and the actions provided allow these reports to structure the surveyed area, to identify similarities not evident from the original description of each individual artifact, to provide design guidance, and even to generalize beyond the scope of the report. Though, it is not uncommon for state-of-the-art reports to develop a customized task categorization for its scope. For example, Alsallakh et al. [4] categorize set visualization techniques using a list of 26 low-level objectives specific to set-typed data. Dealing with visualization systems for electronic

health records, Rind et al. [46] extend the user intent categorization [70] with 20 more concrete subintents and use these for categorization.

Yet, the number of task categorizations can be confusing and in addition the abstract tasks found therein are quite heterogeneous. This is also illustrated in the typology of Brehmer and Munzner [9] that reclassifies abstract tasks from a wide range of existing categorizations. For example, it maps the abstract tasks by Shneiderman [57] to produce, summarize, navigate, filter, and record. Thus, they occur for both why and how as well as at different levels of their typology. Therefore, we propose to use the dimensions of the task cube to classify and compare these categorizations of abstract tasks.

Method. We conducted a non-exhaustive survey of task frameworks in the visualization literature and identified 37 categorizations (Table 1). Then we classified each categorization (1) as objectives versus actions, (2) by three levels of composition, and (3) by four types of abstraction. As we could not identify general classification rules for composition, we made groups of high-level, intermediate-level, and low-level tasks that are as homogeneous as possible. For abstraction, we distinguish between generic abstractions and abstractions that are tailored for a data type, for a domain, or towards the architecture of tools. In addition, we inform about the expressiveness of each categorization by reporting the number of categories, the number of categories in each dimension, or similar information for more complex categorizations. The classification was first proposed by the first author and then discussed and revised by all authors until consensus was reached.

Results. Table 1 provides an overview of the surveyed task categorizations and their classification. Along the *why/how* dichotomy a majority of 24 categorizations is primarily for objectives, while 13 categorizations address actions. For this, we split Brehmer and Munzner’s multi-level typology [9] into its why and how parts. In Roth’s framework [48] we described three categorizations separately: goals and objectives for *why* and operations for *how*. There were a few ambiguities: Schulz et al.’s five-dimensional space [52] has four dimensions relating to *why* (goal, characteristics, target, cardinality) and one for *how* (means). Yi et al. [70] categorize interaction techniques along user intent. Nevertheless their objectives are characterized very similar to actions: e.g., “show me something conditionally” is almost equivalent to a filter action. The taxonomy of Valiati et al. [65] mixes objectives like “identify patterns” with actions like “zoom”.

For composition, we identify only few categorizations for high-level objectives. For example, explore/confirm/present are often found categories or Amar and Stasko [6] present the prototypical tasks of “complex decision making, especially under uncertainty”, “learning a domain”, “identifying the nature of trends”, and “predicting the future”. This is not unexpected as real-world high-level objectives can be very specific to a domain problem. Many objective categorizations can be positioned at intermediate-level, low-level, or both. These ambiguities result from some categorizations having a broad scope (e.g., elementary versus synoptic tasks [7]) and that their abstract objectives can describe for various real-world objectives at these composition levels. Action categorizations are at a mostly low composition level, which is obvious because characteristic sequences of

actions are specific to the concrete situations they emerged from. For this reason, there is also no abstract task frameworks for high-level actions. A few categorizations encompass intermediate-level actions such as ‘overview’ and ‘relate’ [57] or ‘cognitive offloading’ [29].

Among the four types of abstraction, generic abstractions are most common for objectives (12) and actions (8). This can be explained that our survey is not exhaustive of the task categorizations collected for state-of-the-art reports and design studies. Still there is a large proportion of objective categorizations tailored by data type (9) – most frequently networks and/or time.

5. DISCUSSION

We looked at the task cube from various angles of visualization design, evaluation, and research in the previous sections. Reflecting on the dimensions of the task cube, we now observe some general implications.

Usefulness of abstract task frameworks. The long yet non-exhaustive list of task categorizations in Table 1 highlights a large body of research towards theoretical understanding of tasks. Above, we showed various visualization scenarios that require visualization designers, experimenters, and researchers to explicitly consider concrete tasks. But they need abstract task frameworks to generalize such tasks beyond the particular application context: to systematically compare visualization artifacts in a survey, to generate guidelines from evaluation results, and to apply such guidelines and surveys for designing new artifacts. In addition, abstract task lists can be used to check if any relevant tasks still need to be addressed. Our survey shown in Table 1 supports visualization researchers and designers in finding a categorizations suitable for their scenario and application context.

Suitable levels of composition. In capturing the *how* perspective, actions take a pivotal level. Defined as discrete step, they should be at the lowest composition level of which the user is consciously aware. However, this is a matter of context and in particular the user’s experience, which granularity the user regards as discrete steps [37; 44]. One user’s action pattern might be another user’s action and this might also change within a single user over time as he or she progresses from novice to expert. For example, a user might learn to click through menus and dialog boxes unconsciously because he or she frequently needs to change a setting that is hidden there. Then again, visualization designers have also the possibility to detect the need for such action patterns and provide dedicated interaction techniques as shortcuts.

Likewise, objectives can be broken down to increasingly lower levels. However, decomposition into trivially small sub-objectives such as ‘read value 1, read value 2 . . .’ is often not practical and differs from how users realistically solve objectives. Combining human perception and visualization techniques, they can detect patterns at a larger scope. For example, they can spot clusters in a scatter plot or judge the trend in a line plot without consecutively reading the values encoded for individual data records. Therefore, it is important to consider objectives at an adequate level of composition, in particular when evaluating visualization techniques.

Open-endedness of objectives. Visualization is often characterized as undirected exploration or even as casual endeavor leading to serendipitous results. In contrast, the

Table 1: Survey of abstract objective and action categorizations found in theoretical task frameworks, state-of-the-art reports, and domain characterization studies (not exhaustive).

task categorizations	perspective		composition			abstraction				number of categories (in each dimension)
	why	how	HI	IN	LO	GE	DA	DO	TO	
explore/confirm/present [53]	•		•	◦		•				3
Amar and Stasko [6]: prototypical analysis tasks	•		•			•				4
Roth [48]: goals (procure/predict/prescribe)	•		•	◦		•				3
Thomas and Cook [63, p. 35]	•		•			•				3
Amar et al. [5]: low-level components	•			•	•	•				10
Andrienko and Andrienko [7]	•			◦	•	•				2 x 3 x 2
Brehmer and Munzner [9]: why	•			•	•	•				4 x 4 x 3
Munzner [36]: why/actions & targets	•			•	•	•	◦			6 x 4 x 3 x 11
Roth [48]: objectives	•			•	◦	•				5
Wehrend and Lewis [68]	•			•	•	•				11
Ahn et al. [1]: network evolution	•			•	◦		•			3 x 2 x 10
Alsallakh et al. [4]: set-typed data	•				•		•			26
Brehmer et al. [10]: dimensionally-reduced data	•			◦	•		•			6 (maps to [9; 36])
Kerracher et al. [23]: temporal graphs	•			◦	•		•			extends [7]
Lammarsch et al. [27]: time-oriented data	•			◦	•		•			extends [7]
Lee et al. [28]: graphs	•				•		•			11
MacEachren [30, p. 316]: aspects of time	•			•	◦		•			7
Pretorius et al. [45]: multivariate networks	•			•	•		•			25 (maps to [65])
Roth [48]: objectives & operands in cartography	•			•	•		•			5 x 3
Suo [62]: high level tasks in network security	•		•					•		5
Meyer et al. [32]: comparative genomic	•			◦	•			•		14
Schulz et al. [52]: 5-dimensional design space	•	◦		•	•	•	◦			$3 \times \geq 8 \times 4 \times 3 \times \geq 11$
Yi et al. [70]: intents	•	◦		•	•	•				7
Valiati et al. [65]: multidimensional data	•	◦		•	•		•			7
Brehmer and Munzner [9]: how		•			•	•				11
Munzner [36]: how/idioms		•			•	•				11 (different from [9])
Chuah and Roth [11]		•			•	•				18
Gotz and Zhou [17]: actions		•			•	•				21
Liu and Stasko [29]: focus on mental models		•		◦	•	•				6
Sedig and Parsons [54]: action patterns		•			•	•				32
Suo [62]: low-level tasks		•			•	•				32
von Landesberger et al. [66]		•			•	•				3 x 4
Roth [48]: operators on spatial data		•			•		•			17
Rind et al. [46]: electronic health records		•			•			•		20 (extends [70])
Shneiderman [57]		•		◦	•				•	7
Heer and Shneiderman [18]: dynamics		•			•				•	12
Sacha et al. [49]		•			•				•	6

The *why* and *how* columns denote whether the categorization describes objectives or actions. Three composition levels are distinguished: *HI*...high, *IN*...intermediate, and *LO*...low, and four types of abstraction: *GE*...generic, *DA*...data type, *DO*...domain, and *TO*...tool architecture. The symbol • denotes the primary class of the entry; the symbol ◦ represents a partial match, for example if only a few categories fall into this class. The final column describes the expressiveness of the entry by the number of categories or similar details.

term ‘task’ often has connotations as being externally assigned and scheduled within a workflow. We take a pragmatic point of view and understand visualization-supported data analysis as the user steering her or his actions towards solving a reasonably open-ended objective. Such an objective might originate from casual curiosity and/or result in unexpected findings. Still, we assume there is an objective motivating the user’s actions. We also assume that the user has some intellectual interest in the findings. We would not consider an assignment like ‘pan to the year 1983’ or ‘click on the blue rectangle’ as an objective. Such assignments are purely perceptual and mechanical and do not carry the notion of the user raising an inquiry. Furthermore, objectives

in the visualization field involve questions based on *data*.

This intellectual interest in objectives and their many-to-many mapping to action sequences are our motivations to distinguish between the perspectives *why* and *how*. They are also a major difference to the traditional view of tasks in HCI that follows the assumption of a one-to-one mapping between these perspectives [8]. Yet, such an assumption holds only for operational tasks and not for such creative tasks as they are needed for the open-ended objectives addressed by visualization [31; 58].

Summary of discussion. We assume that much of the confusion and criticism on tasks stems from this traditional view of operational tasks. For example, some visualization

experts reportedly [7, p. 148] even “believe that having a defined task is not (or not always) necessary”. Other HCI experts like Preece et al. [44, p. 410] warn that “the idea of a task is useful in system development – as long as it is used with care” because otherwise current action sequences would be reinstated in future systems, which are too rigid to serve the real objectives of users.

However, we agree with Miksch and Aigner [34], Munzner [35], Andrienko and Andrienko [7], and many other visualization experts that tasks are a useful concept for evaluation throughout visualization design and development. Yet, we propose to use the more precise terminology of ‘objectives’ and ‘actions’ instead of the ambiguous term ‘task’. Furthermore, we recommend to understand perspective, level of composition, and level of abstraction as nonredundant dimensions. While these dimensions are clearly distinct in a theoretical view of the conceptual space, it can be difficult to separate them in practice. This is also reflected in a number of ambiguities in our survey of abstract task categorizations. However these ambiguities might also be inherent as a 3D-covariance that deserves deeper investigation.

6. CONCLUSIONS AND FUTURE WORK

In this paper we critically analyzed the usage of the term ‘task’ in visualization and human-computer interaction literature. We propose to use ‘objective’ and ‘action’ as a more suitable terminology that reduces ambiguity and allow visualization researchers to better formulate their contributions. In addition, we identified a three-dimensional conceptual space of user tasks in visualization with abstraction, composition, and perspective as orthogonal dimensions. We looked at objectives and action in a concrete visualization process in various visualization design and evaluation scenarios, and in state-of-the-art surveys and theoretical frameworks. Finally, we emphasized the usefulness of tasks, the importance of choosing adequate composition levels, and focusing on open-ended objectives.

Concluding this article, we can identify two important areas for future research: First, to supply the visualization community with design guidelines [33], we need not only further empirical work but also need to abstract, aggregate and compare these results. Second, tasks and in particular objectives are important to visualization as a objective-driven process. Therefore, visualization artifacts should take users’ intents into account either from explicit input or by auto-detection.

Acknowledgments

We thank the anonymous reviewers, the participants of BELIV 2014, and Bilal Alsallakh for their feedback and suggestions.

Funding

This work was supported by the Austrian Science Fund FWF [grant numbers P25489, P22883]; the Austrian Federal Ministry for Transport, Innovation and Technology [grant number 845598]; and the Austrian Federal Ministry of Science, Research and Economy [grant number 822746].

References

- [1] J. Ahn, C. Plaisant, and B. Shneiderman. A task taxonomy for network evolution analysis. *IEEE Trans. Visualization and Computer Graphics*, 20(3):365–376, 2014. doi:10.1109/TVCG.2013.238.
- [2] W. Aigner. Understanding the role and value of interaction: First steps. In S. Miksch and G. Santucci, editors, *Proc. Int. Workshop Visual Analytics (EuroVA)*, pages 17–20, Goslar, Germany, 2011. Eurographics. doi:10.2312/PE/EuroVAST/EuroVA11/017-020.
- [3] W. Aigner, A. Rind, and S. Hoffmann. Comparative evaluation of an interactive time-series visualization that combines quantitative data with qualitative abstractions. *Computer Graphics Forum*, 31(3):995–1004, 2012. doi:10.1111/j.1467-8659.2012.03092.x.
- [4] B. Alsallakh, L. Micalef, W. Aigner, H. Hauser, S. Miksch, and P. Rodgers. Visualizing sets and set-typed data: State-of-the-art and future challenges. In R. Borgo, R. Maciejewski, and I. Viola, editors, *Eurographics Conf. Visualization (EuroVis) – State of The Art Reports*, 2014. doi:10.2312/eurovisstar.20141170.
- [5] R. Amar, J. Eagan, and J. T. Stasko. Low-level components of analytic activity in information visualization. In *Proc. IEEE Symp. Information Visualization, INFOVIS*, pages 111–117, 2005. doi:10.1109/INFVIS.2005.1532136.
- [6] R. A. Amar and J. T. Stasko. Knowledge precepts for design and evaluation of information visualizations. *IEEE Trans. Visualization and Computer Graphics*, 11(4):432–442, 2005. doi:10.1109/TVCG.2005.63.
- [7] N. Andrienko and G. Andrienko. *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer, Berlin, 2006. doi:10.1007/3-540-31190-4.
- [8] M. Beaudouin-Lafon. Designing interaction, not interfaces. In *Proc. Working Conf. Advanced Visual Interfaces, AVI*, pages 15–22. ACM, 2004. doi:10.1145/989863.989865.
- [9] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Trans. Visualization and Computer Graphics*, 19(12):2376–2385, 2013. doi:10.1109/TVCG.2013.124.
- [10] M. Brehmer, M. Sedlmair, S. Ingram, and T. Munzner. Visualizing dimensionally-reduced data: Interviews with analysts and a characterization of task sequences. In H. Lam, P. Isenberg, T. Isenberg, and M. Sedlmair, editors, *Proc. 5th Workshop Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV*, pages 1–8. ACM, 2014. doi:10.1145/2669557.2669559.
- [11] M. Chuah and S. Roth. On the semantics of interactive visualizations. In *Proc. IEEE Symp. Information Visualization*, pages 29–36, 1996. doi:10.1109/INFVIS.1996.559213.
- [12] K. Cook, G. Grinstein, and M. Whiting. VAST challenge 2014: The Kronos incident, 2014. <http://www.vacommunity.org/VAST+Challenge+2014> (cited 9 Feb 2015).
- [13] A. Cooper, R. Reimann, and D. Cronin. *About Face 3: The Essentials of Interaction Design*. John Wiley & Sons, 2012.

- [14] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering reasoning processes from user interactions. *IEEE Computer Graphics and Applications*, 29(3):52–61, 2009. doi:10.1109/MCG.2009.49.
- [15] C. Forsell and M. Cooper. An introduction and guide to evaluation of visualization techniques through user studies. In W. Huang, editor, *Handbook of Human Centric Visualization*, pages 285–313. Springer, 2014. doi:10.1007/978-1-4614-7485-2_11.
- [16] G. Fuchs. *Task-based Adaptation of Graphical Content in Smart Visual Interfaces*. Ph.D. thesis, Universität Rostock, 2011.
- [17] D. Gotz and M. X. Zhou. Characterizing users’ visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009. doi:10.1057/ivs.2008.31.
- [18] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Queue*, 10(2):30:30–30:55, 2012. doi:10.1145/2133416.2146416.
- [19] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Trans. Visualization and Computer Graphics*, 19(12):2818–2827, 2013. doi:10.1109/TVCG.2013.126.
- [20] T. J. Jankun-Kelly, K.-L. Ma, and M. Gertz. A model and framework for visualization exploration. *IEEE Trans. Visualization and Computer Graphics*, 13(2):357–369, 2007. doi:10.1109/TVCG.2007.28.
- [21] W. Javed, B. McDonnell, and N. Elmqvist. Graphical perception of multiple time series. *IEEE Trans. Visualization and Computer Graphics*, 16(6):927–934, Dec. 2010. doi:10.1109/TVCG.2010.162.
- [22] J. Jo, J. Huh, J. Park, B. Kim, and J. Seo. LiveGantt: Interactively visualizing a large manufacturing schedule. *IEEE Trans. Visualization and Computer Graphics*, 20(12):2329–2338, 2014. doi:10.1109/TVCG.2014.2346454.
- [23] N. Kerracher, J. Kennedy, and K. Chalmers. A task taxonomy for temporal graph visualisation. *IEEE Trans. Visualization and Computer Graphics*, 21(10):1160–1172, 2015. doi:10.1109/TVCG.2015.2424889.
- [24] H. Lam. A framework of interaction costs in information visualization. *IEEE Trans. Visualization and Computer Graphics*, 14(6):1149–1156, 2008. doi:10.1109/TVCG.2008.109.
- [25] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Trans. Visualization and Computer Graphics*, 18(9):1520–1536, 2012. doi:10.1109/TVCG.2011.279.
- [26] T. Lammarsch, W. Aigner, A. Bertone, S. Miksch, and A. Rind. Towards a concept how the structure of time can support the visual analytics process. In S. Miksch and G. Santucci, editors, *Proc. Int. Workshop Visual Analytics, EuroVA, in conjunction with EuroVis 2011*, pages 9–12, Goslar, Germany, 2011. Eurographics. doi:10.2312/PE/EuroVAST/EuroVA11/009-012.
- [27] T. Lammarsch, A. Rind, W. Aigner, and S. Miksch. Developing an extended task framework for exploratory data analysis along the structure of time. In K. Matkovic and G. Santucci, editors, *Proc. Int. EuroVis Workshop on Visual Analytics, EuroVA*, pages 31–35, Goslar, 2012. Eurographics. doi:10.2312/PE/EuroVAST/EuroVA12/031-035.
- [28] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry. Task taxonomy for graph visualization. In E. Bertini, C. Plaisant, and G. Santucci, editors, *Proc. 2006 AVI workshop BEyond time and errors: novel evaluation methods for Information Visualization, BELIV*, pages 81–85. ACM, 2006. doi:10.1145/1168149.1168168.
- [29] Z. Liu and J. T. Stasko. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Trans. Visualization and Computer Graphics*, 16(6):999–1008, 2010. doi:10.1109/TVCG.2010.177.
- [30] A. M. MacEachren. *How Maps Work*. Guilford Press, New York, 1995.
- [31] E. Mayr, M. Smuc, and H. Risku. Many roads lead to Rome: Mapping users’ problem-solving strategies. *Information Visualization*, 10(3):232–247, 2011. doi:10.1177/1473871611415987.
- [32] M. Meyer, T. Munzner, and H. Pfister. MizBee: A multiscale synteny browser. *IEEE Trans. Visualization and Computer Graphics*, 15(6):897–904, 2009. doi:10.1109/TVCG.2009.167.
- [33] M. Meyer, M. Sedlmair, P. S. Quinan, and T. Munzner. The nested blocks and guidelines model. *Information Visualization*, 14(3):234–249, 2015. doi:10.1177/1473871613510429.
- [34] S. Miksch and W. Aigner. A matter of time: Applying a data-users-tasks design triangle to visual analytics of time-oriented data. *Computers & Graphics*, 38:286–290, 2014. doi:10.1016/j.cag.2013.11.002.
- [35] T. Munzner. A nested model for visualization design and validation. *IEEE Trans. Visualization and Computer Graphics*, 15(6):921–928, 2009. doi:10.1109/TVCG.2009.111.
- [36] T. Munzner. *Visualization Analysis and Design*. AK Peters, 2014.
- [37] B. A. Nardi. Studying context: A comparison of activity theory, situated action models, and distributed cognition. In B. A. Nardi, editor, *Context and Consciousness: Activity Theory and Human-Computer Interaction*, pages 69–102. MIT Press, 1996. Preprint retrieved from CiteSeer.
- [38] D. A. Norman. Cognitive engineering. In S. W. Draper and D. A. Norman, editors, *User Centered System Design*, pages 31–61. Lawrence Erlbaum, 1986.
- [39] D. A. Norman. Human-centered design considered harmful. *interactions*, 12(4):14–19, 2005. doi:10.1145/1070960.1070976.
- [40] C. North, R. Chang, A. Endert, W. Dou, R. May, B. Pike, and G. Fink. Analytic provenance: process+interaction+insight. In *Proc. 2011 Ann. Conf. Ext. Abstracts Human Factors in Computing Systems, CHI EA*, pages 33–36. ACM, 2011.

doi:10.1145/1979742.1979570.

- [41] F. Paterno, C. Mancini, and S. Meniconi. Concur-TaskTrees: A diagrammatic notation for specifying task models. In S. Howard, J. Hammond, and G. Lindgaard, editors, *Proc. IFIP TC13 Int. Conf. Human-Computer Interaction, INTERACT*, pages 362–369. Springer, 1997. doi:10.1007/978-0-387-35175-9_58.
- [42] W. A. Pike, J. Stasko, R. Chang, and T. A. O’Connell. The science of interaction. *Information Visualization*, 8(4):263–274, 2009. doi:10.1057/ivs.2009.22.
- [43] M. Pohl, S. Wiltner, S. Miksch, W. Aigner, and A. Rind. Analysing interactivity in information visualisation. *KI – Künstliche Intelligenz*, 26:151–159, 2012. doi:10.1007/s13218-012-0167-6.
- [44] J. Preece, Y. Rogers, H. Sharp, D. Benyon, S. Holland, and T. Carey. *Human-Computer Interaction*. Addison-Wesley, Harlow, UK, 1994.
- [45] J. Pretorius, H. C. Purchase, and J. T. Stasko. Tasks for multivariate network analysis. In A. Kerren, H. C. Purchase, and M. O. Ward, editors, *Multivariate Network Visualization*, pages 77–95. Springer, 2014. doi:10.1007/978-3-319-06793-3_5.
- [46] A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman. Interactive information visualization to explore and query electronic health records. *Foundations and Trends in Human-Computer Interaction*, 5(3):207–298, 2013. doi:10.1561/1100000039.
- [47] A. Rind, W. Aigner, M. Wagner, S. Miksch, and T. Lammarsch. User tasks for evaluation: Untangling the terminology throughout visualization design and development. In H. Lam, P. Isenberg, T. Isenberg, and M. Sedlmair, editors, *Proc. 5th Workshop Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV*, pages 9–15. ACM, 2014. doi:10.1145/2669557.2669568.
- [48] R. E. Roth. An empirically-derived taxonomy of interaction primitives for interactive cartography and geovisualization. *IEEE Trans. Visualization and Computer Graphics*, 19(12):2356–2365, 2013. doi:10.1109/TVCG.2013.130.
- [49] D. Sacha, A. Stoffel, F. Stoffel, K. Kwon, Bum Chul, G. Ellis, and D. A. Keim. Knowledge generation model for visual analytics. *IEEE Trans. Visualization and Computer Graphics*, 20(12):1604–1613, 2014. doi:10.1109/TVCG.2014.2346481.
- [50] P. Saraf, P. Butler, and N. Ramakrishnan. Safeguarding Abila: Discovering evolving activist networks. In *VAST Challenge, Proc. IEEE Conf. Visual Analytics Science and Technology*, 2014.
- [51] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Trans. Visualization and Computer Graphics*, 11(4):443–456, 2005. doi:10.1109/TVCG.2005.53.
- [52] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A design space of visualization tasks. *IEEE Trans. Visualization and Computer Graphics*, 19(12):2366–2375, 2013. doi:10.1109/TVCG.2013.120.
- [53] H. Schumann and W. Müller. *Visualisierung: Grundlagen Und Allgemeine Methoden*. Springer, Berlin, 2000.
- [54] K. Sedig and P. Parsons. Interaction design for complex cognitive activities with visual representations: A pattern-based approach. *AIS Transactions on Human-Computer Interaction*, 5(2):84–133, 2013.
- [55] M. Sedlmair, A. Frank, T. Munzner, and A. Butz. RelEx: Visualization for actively changing overlay network specifications. *IEEE Trans. Visualization and Computer Graphics*, 18(12):2729–2738, 2012. doi:10.1109/TVCG.2012.255.
- [56] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi:10.1109/TVCG.2012.213.
- [57] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. IEEE Symp. Visual Languages, VL*, pages 336–343, 1996. doi:10.1109/VL.1996.545307.
- [58] B. Shneiderman. Creativity support tools: Accelerating discovery and innovation. *Commun. ACM*, 50(12):20–32, 2007. doi:10.1145/1323688.1323689.
- [59] Y. B. Shrinivasan and J. J. van Wijk. Supporting the analytical reasoning process in information visualization. In *Proc. ACM SIGCHI Conf. Human Factors in Computing Systems, CHI*, pages 1237–1246, 2008. doi:10.1145/1357054.1357247.
- [60] M. Smuc, E. Mayr, T. Lammarsch, W. Aigner, S. Miksch, and J. Gärtner. To score or not to score? tripling insights for participatory design. *IEEE Computer Graphics and Applications*, 29(3):29–38, 2009. doi:10.1109/MCG.2009.53.
- [61] D. Sprague and M. Tory. Exploring how and why people use visualizations in casual contexts: Modeling user goals and regulated motivations. *Information Visualization*, 11(2):106–123, 2012. doi:10.1177/1473871611433710.
- [62] X. Suo. *A Task-Centered Visualization Design Environment and a Method for Measuring the Complexity of Visualization Designs*. Computer science dissertations, Georgia State University, Atlanta, GA, 2009.
- [63] J. J. Thomas and K. A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE, 2005.
- [64] C. Tominski, G. Fuchs, and H. Schumann. Task-driven color coding. In *Proc. 12th Int. Conf. Information Visualisation, IV*, pages 373–380. IEEE, 2008. doi:10.1109/IV.2008.24.
- [65] E. R. A. Valiati, M. S. Pimenta, and C. M. D. S. Freitas. A taxonomy of tasks for guiding the evaluation of multidimensional visualizations. In E. Bertini, C. Plaisant, and G. Santucci, editors, *Proc. 2006 AVI workshop BEyond time and errors: novel evaluation methods for Information Visualization, BELIV*. ACM, 2006. doi:10.1145/1168149.1168169.
- [66] T. von Landesberger, S. Fiebig, S. Bremm, A. Kuijper, and D. W. Fellner. Interaction taxonomy for tracking of user actions in visual analytics applications. In W. Huang, editor, *Handbook of Human*

Centric Visualization, pages 653–670. Springer, 2013. doi:[10.1007/978-1-4614-7485-2_26](https://doi.org/10.1007/978-1-4614-7485-2_26).

- [67] M. Wagner, W. Aigner, A. Rind, R. Luh, H. Dornhackl, K. Kadletz, and P. Tavalato. Problem characterization and abstraction for visual analytics in behavior-based malware pattern analysis. In K. Whitley, S. Engle, L. Harrison, F. Fischer, and N. Prigent, editors, *Proc. 11th Workshop Visualization for Cyber Security, VizSec*, pages 9–16. ACM, 2014. doi:[10.1145/2671491.2671498](https://doi.org/10.1145/2671491.2671498).
- [68] S. Wehrend and C. Lewis. A problem-oriented classification of visualization techniques. In *Proc. 1st IEEE Conf. Visualization*, pages 139–143, 1990. doi:[10.1109/VISUAL.1990.146375](https://doi.org/10.1109/VISUAL.1990.146375).
- [69] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: Concepts and methods. In *Proc. ACM SIGCHI Conf. Human Factors in Computing Systems, CHI*, pages 801–810. ACM, 2006. doi:[10.1145/1124772.1124890](https://doi.org/10.1145/1124772.1124890).
- [70] J. S. Yi, Y. A. Kang, J. T. Stasko, and J. A. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans. Visualization and Computer Graphics*, 13(6):1224–1231, 2007. doi:[10.1109/TVCG.2007.70515](https://doi.org/10.1109/TVCG.2007.70515).
- [71] J. Zhao, F. Chevalier, E. Pietriga, and R. Balakrishnan. Exploratory analysis of time-series with ChronoLenses. *IEEE Trans. Visualization and Computer Graphics*, 17(12):2422–2431, 2011. doi:[10.1109/TVCG.2011.195](https://doi.org/10.1109/TVCG.2011.195).