Cognitive Stages in Visual Data Exploration

M. Adil Yalçın
University of Maryland
College Park, MD 20742, USA
adil@keshif.me

Niklas Elmqvist
University of Maryland
College Park, MD 20742, USA
elm@umd.edu

Benjamin B. Bederson
University of Maryland
College Park, MD 20742, USA
bederson@umd.edu

ABSTRACT

Data exploration requires forming analysis goals, planning actions and evaluating results effectively, all of which are complex cognitive activities. Therefore, the data exploration and analysis process can be improved through a principled and comprehensive approach to analyzing the cognitive activities of the user given a data exploration tool. However, many taxonomies and evaluations focus on a specific tool or specific design guides instead of cognitive activities comprehensively. In this paper, we first present the *Cognitive Exploration Framework* that identifies six stages of cognitive activities in visual data exploration. These stages are a combination of two activities—planning and assessing—across data analysis, interaction, and visualization. Cognitive barriers in each stage can lower the success and speed of data exploration. The framework also identifies the factors of decision-making, existing knowledge and motivation that influence cognitive activities. We argue that cognitive stages can be supported by improving the design of tools rather than their computing capabilities. We demonstrate how the framework clarifies the structured relationship between design guides to specific cognitive stages. In particular, the framework can also be used to guide evaluation of data exploration tools. To reveal cognitive barriers in each stage, we focused on the failures instead of success stories, and on motivating self-driven open-ended exploration instead of using benchmarked tasks on fixed datasets. With these goals, we studied short-term casual use of an exploratory tool by novices with limited training. Our results reveal cognitive barriers across all stages. We also discuss directions for future research and applications.

CCS Concepts

• **Human-centered computing~HCI theory, concepts and models**
• Human-centered computing~Empirical studies in visualization

Keywords

Cognition, data analysis, data exploration, evaluation, sensemaking, data visualization, user interfaces.

1. INTRODUCTION

The value of data can be measured by the knowledge we can extract from it. Visual tools support exploration for knowledge discovery by creating an interactive dialogue with data. In this paper, we focus on the role of a *data explorer* with a primary goal of understanding data by developing and answering questions. This is in contrast to consuming pre-extracted knowledge from a presentation (such as a news story), communicating results [22], or designing data exploration spaces/interfaces for other users [5].

Visualization can amplify people's ability to comprehend data [8]. However, using visual tools for data analysis also requires other cognitive activities, such as forming analysis goals and interaction plans. Barriers to effective cognition can lead us to fruitless paths, inaccurate or false knowledge, lost time, or even the abandonment of exploration because of confusion and frustration. Existing work in modeling visualization or cognitive activities in exploration tend to be frameworks that focus on system components [8], [9], [21], empirical results from specific tools and study setups [16], [29], [30] or surveys [31]. Little work has focused on a comprehensive analysis of the cognition on visual data exploration.

In this paper, we present the *Cognitive Exploration Framework* (*CEF*) for visual data exploration, a structured overview of six cognitive stages in data exploration as the combination of planning and assessing activities on data analysis, interaction, and visualization. We identify the factors of decision-making, knowledge and motivation in relation to cognitive activities. By its comprehensive coverage of cognitive activities, the framework can be used to improve and evaluate the design of exploratory tools. First, we demonstrate the rhetorical power of *CEF* by using it to categorize a large number of concrete design guides with respect to stages of cognition. Then, in order to use *CEF* as a lens to evaluate tools, we propose an observational study approach that focuses on identifying failures and challenges in open-ended exploration instead of performance on benchmarked tasks or insights [46]. Our results from evaluation of an exploratory tool with novices in a casual setting showcase the inferential power of *CEF*.

2. RELATED WORK

We summarize the related work on the data-driven sensemaking, cognition, barriers and costs in visualization and interaction.

2.1 Sensemaking and Data Visualization

Sensemaking is an iterative process of gathering and representing information, developing insights through manipulation, and producing knowledge [55]. The information visualization reference model [8], [9] models visualization pipeline from a system point-of-view as transitions between data, analytical abstraction, visualization abstraction, and view. A nested model [38] can be used to evaluate such systems. Yet, these approaches are not based on cognitive processes in visual exploration. Information foraging [42] describes information search behavior using an analogy with animals hunting and gathering food. However, it does not model the data interfaces, interaction, and the analytical process. The data/frame theory of sensemaking [28] argues that sensemaking is composed of cycles of *(i)* elaborating a mental frame, *(ii)* preserving a frame, and *(iii)* reframing. While it models a reasoning process, it does not model the concrete roles of interaction and visualization, and cannot explicitly guide on supporting these processes.

BELIV '16, October 24 2016, Baltimore, MD, USA

DOI: http://dx.doi.org/10.1145/2993901.2993902

2.2 Cognition for Sensemaking

Higher mental processes such as attention, language use, memory, perception, problem solving, and thinking, are the focus of cognitive psychology [14]. Cognition is therefore closely related to sensemaking and data visualization. Card et al. [8] define externalized cognition as the use of an external object to reduce mental effort and memory demands when performing a task. David Kirsh [27] extends the role of external representations into rearrangement, persistence, independence, reformulation, and natural encoding, the use of multiple representations, construction, and simplification of control. In a reverse perspective, Liu and Stasko [35] describe mental models as the internal, structural, behavioral, and functional analogues of external visualization systems. They argue that interaction primarily enables external anchoring, information foraging, and cognitive offloading. *Distributed cognition* models transitions across cognitive representations, and can be applied to infovis [34]. Walny et al. [57] studied data-sketching as an external representation of data understanding. Their analysis focuses on finalized sketches as the artifacts, and not on the cognitive activities explaining *how* the participants created or iterated upon these sketches. While these studies aim to explain the tools as external representations helping cognition, they are primarily explanatory. We aim to close the gap between theory and practice by building a comprehensive and actionable framework, demonstrating its link to design, and its use for evaluating tools.

Shrinivasan [52] presents an analytical reasoning framework with three components, data/knowledge/navigation, which can be supported by special-purpose views in tools. Van Wijk's model of visualization [58] includes perception, knowledge, and exploration as user-level constructs. Green et al. [17] argues these constructs are cognitive processes informing each other. We focus on data exploration using a holistic model covering a wide range of cognitive activities. We identify six cognitive stages, which encompass perception as an assessment activity, and discuss the cognitive influence of knowledge and motivation factors.

2.3 Barriers and Costs in Data Exploration

Generalizing our everyday interactions with the physical world, the gulfs of execution and evaluation [39] is a simple, effective, and widely adopted model. However, it does not fully explain visual data exploration, which involves deep analytical thinking and interaction with abstract data interfaces. Lam [31] presents a framework of seven interaction costs, based on a survey of usability problems reported in 484 papers. Our framework builds upon these works by decoupling cognitive and physical activities, and exclusively focusing on cognition. Amar and Stasko [1] discuss two forms of analytical gaps: (i) worldview gap (what is shown↔what needs to be shown to draw a straightforward representational conclusion) and (ii) rationale gap (perceiving a relationship↔being able to explain the confidence and the usefulness of it). Cognitive stages extend beyond analytical gaps, and we aim to clarify the ambiguous definitions across cognitive activities.

The behavior of novices can reveal barriers that may be reduced or hidden because of existing skills. Grammel et al. [16] conducted an observational study on how novices construct information visualizations. While their study revealed barriers in visualization construction, it did not reflect interactive autonomous data exploration since a mediator created visualizations using verbal descriptions of the participants. Kwon et al. [30] studied behavior of novices to identify visual analytics roadblocks. Participants were given predefined tasks and were offered guidance, which created a partially explorative process and limited the extent of reported roadblocks. Lee et al. [32] identified five cognitive activities in the sensemaking of unfamiliar charts. We argue that the explorer would avoid creating unfamiliar visualizations [16].

Decision-making as a cognitive activity, and its costs and factors, are well-formed within psychology [47]. Yet, decision costs lack a focused discussion in the analytics community. Heer et al. [22] discusses "constraining the parameter space that users have to explore", but only in the context of visualization. Dou et al. [11] studied constrained interactions for solving a math game with empirical results suggesting constraints can increase performance.

3. COGNITIVE EXPLORATION FRAMEWORK

We present the *Cognitive Exploration Framework* (*CEF*) (Figure 1), which identifies six cognitive stages in visual data exploration as a combination of two activities—planning and assessing—across data analysis, visualization, and interaction. Cognitive barriers are impediments that can be observed, categorized, and studied across these orthogonal cognitive stages. In addition, the framework identifies the factors of decision costs, knowledge, and motivation, all of which interact with cognitive stages and influence the exploratory process and outcomes.

3.1 Six Stages of Cognition in Exploration

We describe the cognitive stages using arguments in existing literature below, and show them in exploratory flow in Figure 1.

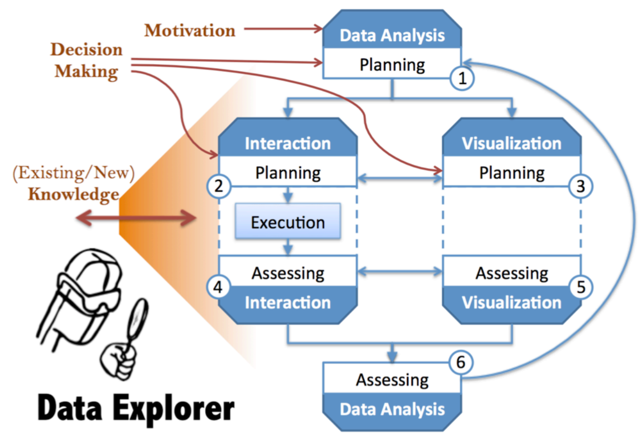

Figure 1. The *Cognitive Exploration Framework* with six stages (shown in blue boxes) and three factors: decision-making, motivation, and existing/new knowledge (shown in red text).

1. Planning Data Analysis: Form goals [29], determine domain parameters [1], characterize task and data [36].

2. Planning Interaction: Form system operations [29], translate queries to attributes [14], execute appropriate interactions [28].

3. Planning Visualization: Design visual mappings / encodings [36] [14], choose appropriate views [28].

4. Assessing Interaction: Evaluate state-change [29], adapt mental model to views [28], the gulf of evaluation [37].

5. Assessing Visualization: Perceive / interpret visualizations [28], visual-cluttering and view-change costs [29].

6. Assessing Data Analysis: Reason about outcomes, observe trends, generate hypotheses, make predictions, assess uncertainty [1], build confidence.

The framework defines *visualization* as the purposefully organized representation of data in an abstract visual language. *Interaction* is the communication between the data and the explorer through the data interface. It encompasses all elements beyond the visual data encoding, such as control panels, buttons, and multiple views. Therefore, our notion of visualization strictly relates to the visual representation of data, and does not cover any interactivity.

In terms of activities in data exploration, *CEF* identifies two activity groups—*planning* and *assessing*—that apply across data analysis, interaction and visualization. Planning activities involve consciously setting goals, making decisions, and identifying courses of individual actions to be taken to reach goals. Assessment activities evaluate the courses of actions taken, data visualizations (through perception), the changes in the interface, and also include reasoning on whether the analytical goals have been answered based on available data, or not. *CEF* models *execution*, such as by mouse or touch, as a physical, *non-cognitive* stage that follows planning interaction, and leads to cognitive assessment stages. It is therefore left out of the scope of cognitive analysis.

In *Cognitive Exploration Framework*, exploration flows from data analysis planning to analytical data assessment to generate knowledge (insights). This is a cyclic and dynamic flow, i.e. exploration can continue with new paths influenced by insights obtained. If a path does not lead to knowledge, or if the explorer is stuck, s/he may retreat to produce new plans or change goals, although time would be lost and motivation may be reduced. The explorer may also act without a purposeful plan, such as selecting a data subset out of curiosity, and reaching insights by observing relations revealed by this actions. Therefore, while the path ideally starts with a well-defined data analysis plan, we recognize it can also be driven by serendipitous interactions.

Next, we discuss factors that influence our model of cognition.

3.2 The Factor of Decision-Making

Increasing options in the exploratory process needs to be assessed not only by what they may enable (richer insights), but also by their cognitive costs. Given many options to choose from, making a decision is harder, and a decision is less likely to be optimal [47]. For example, finding the most effective visualization can be overwhelming given the combination of chart types, glyph types, color, and other visual encodings, especially for novices [16] but also for experienced designers [5]. Avoiding a decision also can be costly. Kobsa reported that Spotfire users tended to use scatterplot, its default visualization, (therefore avoiding chart decision) when another chart type would better fit [29].

CEF generalizes decision costs in data exploration across all planning activities in visualization, interaction, and data analysis. We argue that the options faced in the process of exploration directly influence the decision costs and therefore the cognitive activities. While the examples given above relate to decision factors in visualization, decision-making also applies to data analysis (such as identifying which questions to follow, and which selections to make), and interaction (such as selecting across two alternative actions that may produce the same high-level outcome, or creating a sequence of actions). Every decision can have a positive, or negative, outcome in the exploratory process. *CEF* recognizes and emphasizes the factor of decision making as a potential cost to the cognitive activities in the process of data exploration.

3.3 The Factor of Existing/New Knowledge

The explorer does not only process the data and its interface; s/he also has existing knowledge about the data domain, interface, and visualizations. This knowledge can help across *all* cognitive stages. For example, recalling personal experiences can help forming new queries, and assessing results in a broader context [32]. As the explorer gains more skills, the plans and assessments can improve. However, existing knowledge is limited, non-universal, and varying across people. In addition, knowledge is dynamic, i.e. there is learning during exploration and use of the tool. The explorer iteratively uses, builds, and evaluates knowledge constructs [28]. S/he does not only learn about the explored data, but also about the interface, interactions, and visualizations, which can lead to more effective use of the tool over time.

3.4 The Factor of Motivation

What are the driving forces of the explorer to engage in data exploration? *CEF* identifies potential answers as the *motivation* factor. Motivation can follow the *curiosity*, such as to understand the data content and features. *Being in the flow* is another motivational construct. The flow—the balance between the challenge of a task and user skills—can apply within the context of interface use [4] and visual analysis [17]. *Creativity* is also motivating, and is applicable to data analysis (finding goals), interaction (combining features of the interface), and visualization (finding new forms to see new data perspectives). *Emotions* can also be motivating. Harrison et al. [20] found that emotion (affect) priming can influence perception of visualization. We propose that this result can apply to a wider range of activities in data exploration. Positive mood can increase motivation, and therefore exploration success.

4. DESIGNING FOR COGNITIVE STAGES

In visual data exploration, the data interface becomes the communicative channel between the cognition (mind) and the data. Supporting cognition (and reducing barriers) is therefore most related to the design of the tool interfaces rather than their computational models. In turn, what is the relation between design and the cognitive stages? How can the cognitive barriers be reduced by principled design? To answer these questions, we contribute a new categorization of 29 concrete and common design guides by linking them across six orthogonal stages of the *Cognitive Exploration Framework*. This section can be used to guide and improve the design of data exploration tools. The wide range of principles covered supports the rhetorical power of the *CEF*, which creates an orthogonal space for analyzing cognitive activities.

Our selection of guides is based on the existing practices and literature. Although we aim to present a wide coverage and effective exemplars for each stage, offering a complete list of guides is impossible, and an extensive list is out of our scope. These guides should not be taken as *rules* of design, but rather directions to consider in designing tools that better support cognitive activities.

4.1 Guides for Planning Data Analysis

- Promote overview-to-detail exploration [50]. Starting with the data overview helps the explorer build a high-level mental model. Reveal detailed relations by interaction progressively.

- Show only relevant exploratory paths. Promote never-ending exploration [12]. Prevent queries leading to zero results [18]. Systematic yet flexible discovery [41] enumerates exploratory paths to suggest unexplored areas and communicate progress.

- Make exploration steps easily reversible [12]. This motivates action and reduces decision costs.

- Provide traces of exploration paths. To form new goals, the explorer may use action histories [23].

4.2 Guides for Planning Interaction

- Use direct manipulation [12], [51]. This reduces the cognitive distance between planning and execution through a continuous representation or metaphors of objects in the interface.
- Integrate interface with visualizations [12], [17]. This promotes visual coherence in a single immersive environment. Scented widgets [59] suggest designs on merging visualizations with interface elements such as dynamic query widgets [49]. Legends can also be designed as interactive widgets [43].
- Show only relevant interaction options. Design to provide context; reveal interactions relevant to the selected object. Design based on the context; reveal contextual interfaces only when the explorer interacts with relevant object (e.g. show actions icons on mouse-over).
- Indicate affordances of visual objects clearly [12]. Use visual cues to suggest interactivity [6].
- Design to fit the cognitive and conceptual model of the explorer. Allow searching for concrete data values, expose context of data attributes and their semantic relations, and support partial specification of exploration paths [16].
- Make every step useful and pleasing [12]. An action should not lead to a confusing, ineffective interface.

4.3 Guides for Planning Visualization

The primary means to support cognition in planning visualization is reducing the visualization parameters and options, starting with showing sensible defaults [23].

- Show only appropriate visualization options for the underlying data types and intended tasks. Recommendations may be a short list of suggestions, such as Tableau's "show me" feature [36], which uses a rule-based design on selected attribute types, or a fully automated approach [45]. The context of use can also be considered [15].
- Support alternative visualizations to reveal relations that cannot be explored with existing views. Alternatives should be functional and add minimum decision costs. For example, given cities and their populations, an ordered list would reveal the cities with most/least populations, a histogram would reveal the population distribution, a map would reveal the spatial context, and a line chart would reveal temporal changes.

A common practice in visualization design is *templating*, in which the explorer selects a chart type first, and then decides which attributes to map to template parameters: axes, color, size, and so on. However, using visualization templates can impede cognitive activities because they require the explorer to understand the template parameters to make effective mappings [16]. Thinking is restructured to the terms of the template parameters from the terms of exploratory goals, potentially creating a mismatch of mental representations. Templates can be richer than fixed chart types such as flexible shelf-based systems [54] that construct a parameterized visualization space. We argue that revealing systematic parameters of a visualization design space should not be the basis of constructing visualizations for exploration.

4.4 Guides for Assessing Interaction

- Make system status clearly visible [39].
- Link multiple views on interaction [44]. Having multiple views increases the cognitive load with more visual information to digest. Linking views reveals relations between data representations, and can improve mental models. Linking should be consistent and intuitive.
- Provide real-time feedback after interaction [12]. A visual feedback delay, as short as 500ms, can decrease exploration activity and data coverage [33].
- Animate transitions between interface states [12]. Avoid abrupt changes and provide a sense of direction.

4.5 Guides for Assessing Visualization

- Use effective visual encodings. Graphical perception studies [10] report how accurately and rapidly we perceive data graphics across different encodings.
- Use appropriate scales, grids, labels, legends [22].
- Aim to reduce visual complexity. Avoid overlapping glyphs since they are a basic form of visual complexity.
- Avoid duplicate representations. Duplication of the same data point may increase cognitive efforts, as it requires understanding relations across multiple glyphs of the same data. Each additional glyph also takes screen space, which is a limited resource that should be carefully used.
- Aggregate data, when it cannot effectively fit in limited screen space, and to provide overviews.
- Show conceptual data domain. For example, use matching icons (as glyphs or isographs) and matching colors for categories [48] where appropriate. Show uncertainty [1] when data has an uncertainty measure.
- Animate transitions of data glyphs [24].
- Use available screen space effectively. Adapt the visualizations based on display size.

4.6 Guides for Assessing Data Analysis

- Provide multiple views (perspectives) of data [17], [44]. One visual representation cannot show all aspects of rich data. Simultaneously observing multiple views can reveal relations across individual views.
- Show the semantic context of data [16], such as description of data attributes, categories, and data values.
- Provide analytical models for statistical analysis. The tool can support the explorer to accurately evaluate findings using statistical methods such as hypothesis testing with significance [1].

4.7 Generalized Guides across Stages

- Aim for consistency. Inconsistencies in visualization, interaction, or interface design make it harder to form goals and action sequences, make decisions, perceive data, and the interface state. Therefore, consistency can influence both planning and assessing stages across multiple artifacts.
- Aim for minimalism. Make design as little as possible [37], [56], [60]. Showing only relevant paths and options in context of active state is a form of minimalism that can support planning cognitive activities. Minimalism can also reduce complex systems to fewer components that are easier to evaluate, thus supporting cognition for assessment.

5. EVALUATION TO DETECT BARRIERS

The success of data exploration depends on cognitive activities, and the cognitive barriers therein. We propose an evaluation to better understand the cognitive activities of the analyst/explorer. Specifically, we focus on challenges in data exploration and the cognitive barriers to collect and analyze potentially actionable observations. In this section, we discuss how cognitive activities can be observed per each stage in evaluating an exploration tool, and how the *CEF* provides a high-level structure to the analysis. Our goal is not to describe evaluation of a specific design guide, or a single stage of cognition, such as visualization perception, which require different setups. We don't aim to present new guidelines for design, or a comprehensive analysis of an existing tool. Rather, we present a new evaluation approach as a lens that focuses on and reveals barriers to cognitive activities.

We argue that detecting cognitive barriers requires focusing on failures, such as lack of goals, not being-in-flow, ineffective plans, and invalid insights. This is in contrast to the common practice of searching for success stories of our tools. Using benchmark tasks on fixed datasets does not facilitate autonomous, self-driven exploration. Furthermore, it may fail to motivate participants with a wide range of interests and background, or alienate them. We suggest that the participants should express their interests in selecting data domain and their exploration goals in order to improve their motivation and success. Furthermore, to expose all cognitive activities clearly, participants should be encouraged to interact with the tool directly without guidance by the facilitator. While usability studies commonly focus on physical execution problems and surface-level software use activities on pre-defined, benchmark tasks, our goal extends to reveal the cognitive processes of the user in a natural setting. To summarize, our study protocol aims to position participants as *explorers* aiming to discover meaningful data-driven knowledge in an open-ended setting to answer their own questions based on their interests.

Revealing cognitive activities in depth requires moving beyond basic observations. For example, the *explorer* may want to sort a list alphabetically, interact with various interface components to find this feature, and then give up and change her goal. Detecting such a process as a negative outcome is instrumental to understanding cognitive activities, especially when such tasks cannot be exhaustively enumerated. How can such failed actions and goals

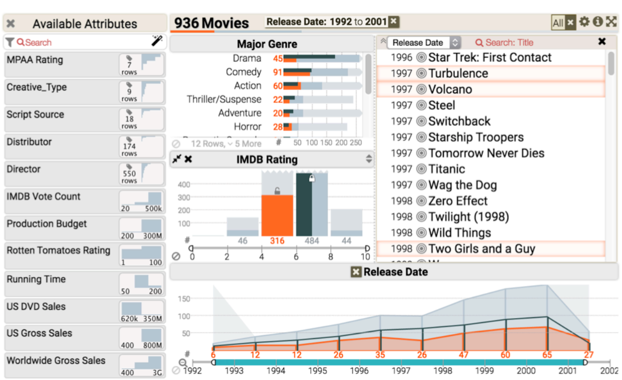

Figure 2. A view from exploring movies. Movies are filtered by release date 1992-2001. Movie names are ordered by release date. The black color in visualizations shows the distribution of movies with IMDB Rating 6-8 (locked selection), and orange shows the range 4-6 (highlight selection with mouse-over), including three highlights on movie list. Left-most panel shows other attributes that can be inserted to the browser.

be observed by the analyst or some algorithm? Software logs [19], eye tracking [53], and brain scans [2] have some, yet limited, power in describing reasoning and exploration processes. Alternatively, encouraging verbal communication and analyzing the discourse can allow observing parts of the cognitive processes [13].

As the basis of our protocol, we suggest that cognitive activities can be revealed with the facilitator observing the exploration process for potential challenges, asking for clarifications, prompting for more communication based on exploratory stages and reasoning behind actions of the participant. These interventions should be minimal and focused on cognitive activities, not a test of knowledge or a measure of success. Surveys and others forms of external cognition can also facilitate communication of cognitive processes. Our position is that, taken together, observations, interventions, surveys and external cognitive methods can lead to identification of a rich set of cognitive activities in data exploration.

With these goals and background, we evaluated a data exploration tool in the open-ended, self-driven exploration protocol with limited training, and analyzed self-reported feedback and cognitive activities using *CEF*. Our study with small number of non-expert participants exemplifies a range of barriers across cognitive stages, suggesting that cognitive stages can be robustly observed using the proposed study protocol and with *CEF* guiding the analysis.

5.1 Keshif: A Data Exploration Tool

In this section, we describe the studied tool for visual data exploration, called Keshif (www.keshif.me) (Figure 2). Its expressiveness focuses on understanding data distributions across fully linked selections in univariate visual summaries and record list. Its visualizations are fixed per data type (categorical, numerical, temporal) with interaction driven by exploratory tasks. Next, we demonstrate the design of the studied tool.

Sally is interested in movies, and finds a tabular movies dataset. Upon loading it, the tool shows that the dataset has 3.2k movies with 16 attributes. In this initial *highest-level overview* with the attribute list and the empty browser, Sally learns attribute names and the distributions of movies within each attribute using its visualization snippet. With her interest in IMDB Ratings, she double-clicks to expand it to a summary in the browser with larger visualization. The tool creates histogram bins on this numeric data based on the summary width and the range of values, aiming to use the space effectively. IMDB Rating disappears from the left panel to have only a single instance of an attribute, avoiding duplication. She then inserts Worldwide Gross Sales to the browser, placed automatically below the rating without overlaps. Distribution of data in this attribute is highly skewed (few movies have very high sales), so the tool applies log transformation for binning, auto-adjusting visualization to the data. Sally then inserts movie titles to the browser. As a unique attribute, the tool lists the titles in the middle panel, with sorting, text-search and ranking features. The record view panel is unique, and each record is only represented once, avoiding duplication by design. She then inserts the genre and the script type to the browser; both are categorical attributes automatically placed to the left panel. When she double-clicks Release Date, a timestamp attribute, the tool visualizes distribution as an area line-chart instead of bars, since connected lines can improve perception of temporal trends. Release Date summary is placed to the bottom panel to widen the chart area for binning on time ranges. At any point, she can (re)arrange summaries across left, right, middle, or bottom panels by drag-and-drop, a direct manipulation design.

Sally then wants to focus on details on data subsets by selection. The tool offers three selection modalities for three goals, visualized using different colors consistently: (i) mouse-over to rapidly highlight selections, (ii) lock a highlighted selection to compare it with other selections, and (iii) click to filter the dataset. When the data is filtered, the total distribution is shown by color, and the filtering state is shown using breadcrumb pattern. With a tight integration of interface and visualizations, summaries support direct interaction on familiar visual forms of labels, numeric counts, and univariate visualizations. Multiple summaries and the record view provide multiple views into the dataset by grouping and listing movies. Bidirectional linking across all types of selections creates consistency, and reveals relations across attributes rapidly. The tool also features part-of scale mode, which enlarges the visual glyphs to their full extent (100%) and enables observing part-of-whole distributions of highlighted or compared selections across record groups.

Filtering and dynamic updates are also designed based on the data type. Categorical summaries support *and/or/not* filtering within, and include text-search when there are more than 20 categories. Numeric summaries support range queries and zooming and panning along the value (horizontal) axis. Upon filtering, categorical attributes are re-ordered with larger measurements on top, unless the attribute is ordinal. To zoom into a numeric range, Sally needs to first filter a range, which makes the zoom-in icon visible (when relevant), and then she can zoom into the filtered range, which changes the icon to zoom-out. Thus, zooming is controlled by only a single contextual button. Upon zooming, the numeric bins are re-adjusted to fit the zoomed range. The tool also animates changes in visualizations and panel layout.

5.2 Study Design

To detect cognitive activities and barriers in exploration, we designed a casual setting with a 15-minute exploration per dataset, and 5-minute training for using the tool. As existing knowledge and extensive training can reduce the barriers that the evaluation aims to detect, we aimed to recruit novices in data analysis, and offered limited training. The participants chose two multivariate, tabular datasets they would like to explore given five options: movies, traffic accidents, passengers of the Titanic, Lego sets, and foodborne disease outbreaks. The record (row) count ranged from 3.2k to 30k, and the attribute (column) count ranged from 8 to 16.

To encourage communication on exploration and emotional states, we implemented an external strategy using printed cards. One group of cards described exploratory process: (i) "*I am trying to find a question.*" (planning data analysis) (ii) "*I am trying to answer a question.*" (planning interaction & visualization) (iii) "*I have an insight.*" (assessing data analysis). Another group of cards focused on negative emotions: "*I feel...*" (i) confused, (ii) undecided, (iii) lost, (iv) bored, and (v) frustrated. The use of cards was not mandatory; the participants could talk on their observations and challenges without picking or pointing to cards.

Procedures and data collection. At the beginning of the study, the participants completed a background survey on demographics (age, sex), existing knowledge in data analysis, visualization, and computer use/interaction, and overall motivation in data exploration, using a talk-aloud protocol. Then, they were trained with a 5-minute video tutorial, which described the tool features while demonstrating data analysis, and 20-slide printout for future reference. After the training video, the facilitator presented the cards, and asked the participants to think aloud while exploring data, and use the cards if appropriate. To gain familiarity with the tool and the study process, the participants explored the training dataset for 5 minutes. Then, they explored two datasets of their interest, 15 minutes each. The facilitator answered questions about the tool based on the training material. While we encouraged self-driven exploration without external tasks, the participants could pick among five sample questions per dataset. After each dataset, the participants completed a survey that encouraged recalling both positive and negative experiences, using ten Likert-scale questions based on [40]. We recorded the screen and the audio in the room during participation in our study. To detect the cognitive barriers, the lead author watched the videos and took note of the problems faced by the participants, and their relevant verbal feedback, including feedback based on the surveys. He then classified them across the six cognitive stages.

Participants. We recruited participants using public message boards. Our participants were non-experts in data visualization and analysis. Our study included pilot-sessions with two participants and reported-sessions with three participants (P1, P2, P3). P1 was a male student in biology, age 18- 24. P2 was a female professional in finance, age 40-49. P3 was a female student in food science, age 18-24. All participants were familiar with basic chart types (bar-charts, histograms, line-charts, maps), and none were familiar with advanced chart types (scatterplots, treemaps, node-link diagrams and ||-coords) by name. The self-reported computer skills were novice (P1, P3), intermediate (P2), and none advanced. All participants had experience with Excel, including basic visualizations, data entry (P1), formulations (P2), and none had experience with other data tools. Their motivation to join the study was curiosity (P1, P2, P3), and earning money (P2); $10 for their 1-hour participation. While this reflects the demographics of the study location, a university campus, their data analysis experience were none (P1) or infrequent (P3), only P2 noting to frequently analyze data "*to figure out the yield on investments.*" The participants were interested in the following domains: movies (P2, P3), traffic accidents (P1), foodborne outbreaks (P1, P3) and *Titanic* passengers (P2). Per each participant, the use of sample questions to bootstrap exploration was: P2-none, P3-1 question, and P3-multiple questions.

5.3 Exemplar Barriers per Cognitive Stage

In this section, we demonstrate the application of the *Cognitive Exploration Framework* for tool evaluation using the proposed protocol. We report exemplar barriers faced by our participants.

5.3.1 Barriers in planning data analysis

Talking about his experience, P1 noted, "*Maybe I felt like I had too much control, but I wasn't ready for it*", and added "*I wasn't quite able to figure out what I wanted to figure out.*" He stated he was overwhelmed at points (by multiple views), noting, "*It's just a lot to take in. A lot of different elements to consider… I don't understand how to put (a lot of information) together.*" P2 set some serendipitous goals, "*Let me see (filter) Clint Eastwood and see what happens.*" When picking sample questions, P1 noted on his motivation, "*I want to find something… that I'd personally want to get the answer to.*" In addition, to save the limited time, P1 did not want to pick questions that looked complicated to answer. Goals were also constrained by the content of data. P3 said, "*(the data) doesn't have enough criteria to give you a definite answer*", as she wanted to relate diseases from fish consumption to fish production per state. To address the information overload, the tool can be designed to offer simplified authoring interfaces, or to encourage step-by-step guided exploration. Sample goals can be provided from simple to complex as the user gets familiarity using the tool.

5.3.2 Barriers in planning interaction

After getting stuck in a question, P1 noted, "*The computer doesn't really know the question that I have (…) I am confused about how to go by answering that question, or if the method I'm using is actually the right way.*" P3 was confused after an ineffective sequence of actions—filtering, locking, and selecting the same histogram bin—where she noted, "*I don't know what exactly I'm trying to do.*" Participants also updated interaction plans and goals given the design and limitations of the tool. To search for specific values, P2 first wanted to alphabetically sort categories and records (not supported), then she used text search, an appropriate strategy. When P2 wanted to sort few movies by year, which could be achieved using sorting dropdown, she hovered the cursor over movies to automatically highlight their year within summaries. Being satisfied with this approach, she discarded her original sorting plan. We also observed some learning challenges with contextual interfaces. P3 wanted to resort categories in reverse, however was not able to easily find the sorting button because it was hidden by default, and shown only on mouse-over in categories. She later suggested, "*If I had more practice with this, I would definitely be in more control.*"

To address the change-of-plan during sorting goal, we updated the design of the tool to include a sorting button within the summary in addition to the sorting option combobox. The tool can also be improved to identify repeated actions to reason about user intent, and suggest relevant actions to help the user plan for interaction.

5.3.3 Barriers in planning visualization

With the selected tool, activities related to planning visualization include aggregate selection modes (highlight, compare, filter) and part-of/absolute mode. This contrasts to the charting tools that would require more careful planning to construct effective visualizations. Therefore, barriers in this stage were not frequently observed. In trying to find the most common food outbreak in different months, P3 filtered through multiple months, while highlighting would be more effective. Another barrier was that participants could not plan to execute part-of scale mode change, as no participant in our study used part-of scale. This may reflect that their questions may not have required such views, but also suggests that the limited knowledge about how this mode could be used effectively. The tool design may be improved to communicate and clarify the use of part-of scale mode to answer related questions.

5.3.4 Barriers in assessing interaction

Failing to consider filtering selections correctly was a common barrier leading to false conclusions about general, or targeted, populations. After unfiltering a selection, P1 said, "*I forgot that I had still filtered everything for the norovirus.*" When P2 wanted to analyze *survivors* of the *Titanic*, she highlighted *non-survivors* and reached a wrong conclusion about their ages. She realized and corrected her mistake shortly after. P3 interpreted the full bar length in a filtered summary to support her misunderstanding that the complete dataset was selected. P3 also misinterpreted how selections are linked across summaries, saying "*If I lock (this bar), there's no way I could compare to (another summary) because they are two different things.*" Overall, tracking multiple selection states was found to be a non-trivial task for the novices in our experiment. The tool can be updated to offer simplified interactivity to reduce confusion on dynamic selection changes.

5.3.5 Barriers in assessing visualization

P1 was confused about what the numbers represent upon selection, saying "*Is this number representing fatal accidents, or just accidents or is it drunk vs. non-drunk. . . . Ok, I didn't realize there are two different colors.*" P2 tried to understand linked highlighting selections by hovering on different bars, observing numbers, and making connections. P3 had trouble observing exact filtering range within the line chart because of its design. The rounding of histogram end-points also lead to wrong interpretations. With maximum duration of movies at 157 minutes, the high end-point of histogram was rounded to 300 minutes, an anomaly of the log scale used by the tool. With this view, P3 interpreted there were movies up to 300 minutes. Real maximum value could be observed by sorting movies in decreasing duration. We later improved the design of our tool by placing the maximum-tip on the scale to the real maximum value, instead of the maximum of the histogram bin range that may exceed true maximum. Filtering range limits can be more explicitly noted in interval summaries, as well as information about what each number presents in the charts under different selection configurations.

5.3.6 Barriers in assessing data analysis

Understanding data semantics was a common challenge. P2 asked "*How do I find the definition of vote count?*", and later removed this summary from the browser. P3 asked, "*What is 'ethnic style, unspecified' (as food type)? That could be anything.*" and then noted, "*This doesn't really affect the program, it's just the data itself.*" Notice that these comments to not reflect to either visualization or the interaction design, and relates to data concepts related to analysis. Unexpected findings raised suspicions, with participants concluding, "*if I'm interpreting right (P1)*", and "*if I'm reading right (P2)*". Acknowledging an inappropriate strategy to reach answers, P1 said, "*I am merely associating these numbers with the question that I have.*" When only 10-20 outbreaks were selected after filtering, P3 concluded about statistical trends and did not discuss limitations of their significance. No participant recognized that some summaries did not include all records, e.g. there were movies without rating information. Another issue was potentially misleading inferences across summaries. When the filtered movies had high-ratings, and kids movies were common, P3 (incorrectly) inferred that kids movies had high ratings, without querying further to confirm her intuition.

To address assessment challenges in data analysis, providing contextual information about metadata would be helpful. Warnings can be presented when few records remain to make statistical conclusions, or missing records can be highlighted explicitly.

5.3.7 The Factor of Existing/New Knowledge

Our participants were non-experts in visual data analytics. We further limited training and studied a casual short-term use to limit the factor of knowledge. We observed this approach influenced the experience and feedback of our participants. P1 said "*It's been a while since I looked at charts… You have to re-familiarize yourself with all the information it represents.*" P2 "*felt discouraged, just in the very beginning, as I was getting used to the tool.*" P3 added "*You never really learn it until you actually try to do it.*" These feedbacks point to the active learning experience of the participants during the use of the tool.

6. DISCUSSION

6.1 Construction of the Framework

We presented the *Cognitive Exploration Framework* to provide a comprehensive overview of cognitive activities, the role of design in cognition, and how barriers to cognitive activities can be the focus of evaluation of tools. To construct the framework, we iteratively identified and refined various arguments about cognition and barriers in related literature (see Section 2) as well as our own experiences in evaluation and interface design. For example, the gulf of execution and evaluation [39] models physical or lower-level cognitive activities, while Lam [31] focuses on interaction-related usability problems. Both models are used to build our framework after separating physical execution stages. Our framework is further enriched and supported by other arguments such as positioning of analytical gaps and activities [1], and results from empirical user studies [5], [16], [30]. Overall, we noticed similar themes across taxonomies and empirical studies stated in different perspectives. We hope that the six-stage orthogonal overview of the Cognitive Exploration Framework and its relation to design and evaluation will provide a concrete, lean basis to understand cognitive challenges of visual data exploration and analysis.

6.2 Implications for Design Guides

Our overview of design guides suggests that existing literature provides many guidelines and discussions for interaction and visualization design. However, high-level data analysis and planning stand as cognitive activities with opportunities for more results and guidelines with future studies. One of the challenges is identifying *how* people reason about data and plan for data analysis. Another challenge is evaluating high-level outcomes of exploration and cognitive planning activities. Equipped with better models for cognition and evaluation methods that expose new metrics and processes, new improvements and guides may be enabled. The results and examples from our user evaluation support that high-level cognitive activities can be analyzed qualitatively by observing failures in user behavior and verbal feedback.

6.3 Evaluating Tools for Cognitive Barriers

To detect cognitive barriers, we designed an experiment with an open-ended exploratory setting, allowed the participants to choose a dataset and exploratory goals of their interest to increase motivation, and applied brief interruptions to encourage the participants to communicate their exploratory process and their negative emotions/experiences in a safe environment. While insight-based methodologies [46] focus on the success stories to quantify the observed value of a tool, a principled way to understand *failures* reveal opportunities for improvement. Our evaluation is a reflection of the open-ended data exploration approach, aiming for the unknown and the intangible in the process of exploratory cognition and generating qualitative, rather than quantitative, value. We showed that *CEF* can be applied in practice to detect and categorize observed barriers on cognition effectively, although we did not create *CEF* on empirical results from this particular study.

Our study design can be replicated or modified to study cognition in more depth. While we used think-aloud protocol and discourse analysis along with actions observed in video and notes taken by the facilitator, this approach has its own limitations, especially for comprehensiveness. This qualitative analysis can be coupled with other forms of behavior tracking, such as software logs and eye movements, to add quantitative support for detecting cognitive activities. Using pair analytics protocol [3], the cognitive stages can be distributed across subject matter expert (high-level cognition in data analysis) and visual analytics expert (low-level cognition in interaction and visualization).

In retrospect, we observed that the participants rarely used the cards to express their emotional and exploration state. While external anchoring may be beneficial to reveal more activities, the participants were either immersed in their data exploration, or was not paying attention to the cards that were displayed on the table next to the study laptop. Embedding these feedback mechanisms on the interface of the tool may make them more prominent. The benefit of such external mechanisms can be studied further to detect if they lead to more communication. As we had a small number of participants, we used the survey as a way to collect more feedback from the participants rather than to build a semi-quantitative analysis. Selected quotes we reported include feedback during the completion of post-exploration survey. We suggest the use of surveys to create opportunities to gather more feedback about the experience of the participants.

Since our goal is to find exemplar barriers in this preliminary study, we did not fully transcribe the sessions that require higher effort and resources. Having more participants, full transcriptions and multiple passes over the recorded material may reveal more cognitive activities in the use of a studied tool.

6.4 Effort Differences across Cognitive Stages

Do all cognitive stages require the same mental effort? Daniel Kahneman [25] argues that our cognitive activities are two-folded: system-1 (thinking fast) and system-2 (thinking slow). System-1 is how we make quick decisions, take short-cuts, apply our cognitive biases, etc. It is less deliberate and more spontaneous. System-2 is how we engage in a more effortful thinking, be more analytical, evaluate facts, and even actions of system-1. We argue that the stages of planning and assessing *data analysis* requires higher cognitive efforts as a slow thinking activity, and that fast thinking activities include perception of visualizations, evaluation of interface and planning for low-level actions respectively. Future research may investigate the differences of effort in cognitive activities in data exploration under various settings.

6.5 Contextual Limitations of the Framework

The framework is developed for self-motivated data exploration of an individual given a specific dataset and a tool that enables exploration. We clarify the limitations of our context below.

Finding data: The framework assumes data has a well-defined structure and content. It does not discuss finding or wrangling data [26]. Such activities, before data planning, can be modeled a recursive exploration for data.

External objectives: If questions are known in advance, exploration is not necessary, just methods to query. Understanding the objective is a natural challenge given predefined objectives [30]. The framework separates external objectives, and values exploration as autonomously evolving goals in a dynamic environment.

Collaboration: We modeled the *explorer* as an individual. Interaction with other explorers brings new cognitive activities, such as creating and following shared plans, and collaborative learning. Collaborative exploration is a richer process beyond our scope.

Tasks: We argue that tasks in exploration originate from cognitive planning activities with discovery purposes. However, *CEF* does not aim to provide a model of tasks (e.g. [7]).

Knowledge representation: The framework does not describe the form of cognition, such as frames, schemata, or propositions [28]. The activities, such as making decisions, evaluating results, and generating knowledge, are transitions on such forms. To guide the explorer through externalized representations, the tool may facilitate capturing extracted knowledge or exploration paths.

7. CONCLUSION

In this paper, we focused on the cognitive activities in open-ended visual data exploration. By identifying planning and assessment as cognitive activities across data analysis, interaction, and visualization, we presented the *Cognitive Exploration Framework* for visual data exploration that is composed of six orthogonal cognitive stages. We used the framework to identify how established design guides interact with the cognitive stages. We then demonstrated application of the framework in evaluating a data exploration tool by focusing on the failures and challenges on self-motivated, autonomous data exploration using data analytics novices with limited training in a casual setting.

While our analysis exemplify a range of barriers tied to the framework (some of which are potentially addressable by incremental design improvements), it also raises questions about how to better support analytical goal formation and analytical evaluations by design. Providing a sense of control and ease while increasing the expressive power of tools is a major challenge. To move beyond the casual setting of our demonstrative user evaluation and to observe complex activities, future studies may increase training, motivation, domain knowledge and skills of the participants. Identifying the influence across cognitive stages and quantifying the differences in efforts can further guide better design of our tools, allowing us to explore data in depth more rapidly.

ACKNOWLEDGMENTS

This work was partially supported by the National Socio-Environmental Synthesis Center (SESYNC) through a grant from the U.S. National Science Foundation (# DBI-1052875) to the University of Maryland. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We wish to thank Ben Shneiderman, Nick Magliocca, and the reviewers for their feedback and discussions.

REFERENCES

[1] R. A. Amar and J. T. Stasko, “Knowledge precepts for design and evaluation of information visualizations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 4, pp. 432–442, Jul. 2005.

[2] E. W. Anderson, K. C. Potter, L. E. Matzen, J. F. Shepherd, G. A. Preston, and C. T. Silva, “A User Study of Visualization Effectiveness Using EEG and Cognitive Load,” *Computer Graphics Forum*, vol. 30, no. 3, pp. 791–800, 2011.

[3] R. Arias-Hernandez, L. T. Kaastra, T. M. Green, and B. Fisher, “Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics,” in *44th Hawaii International Conference on System Sciences (HICSS)*, 2011, pp. 1–10.

[4] B. B. Bederson, “Interfaces for Staying in the Flow,” *Ubiquity*, no. September, pp. 1–1, 2004.

[5] A. Bigelow, S. Drucker, D. Fisher, and M. Meyer, “Reflections on How Designers Design with Data,” in *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*, New York, NY, USA, 2014, pp. 17–24.

[6] J. Boy, L. Eveillard, F. Detienne, and J. Fekete, “Suggested Interactivity: Seeking Perceived Affordances for Information Visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 639–648, Jan. 2016.

[7] M. Brehmer and T. Munzner, “A Multi-Level Typology of Abstract Visualization Tasks,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2376–2385, Dec. 2013.

[8] S. K. Card, J. D. Mackinlay, and B. Shneiderman, “Readings in Information Visualization: Using Vision to Think,” San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999, pp. 579–581.

[9] E. H.-H. Chi and J. T. Riedl, “An operator interaction framework for visualization systems,” in *IEEE Symposium on Information Visualization*, 1998, pp. 63–70.

[10] W. S. Cleveland and R. McGill, “Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods,” *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 531–554, Sep. 1984.

[11] W. Dou, C. Ziemkiewicz, L. Harrison, D. H. Jeong, W. Ribarsky, X. Wang, and R. Chang, “Toward a Deeper Understanding of the Relationship Between Interaction Constraints and Visual Isomorphs,” *Information Visualization*, vol. 11, no. 3, pp. 222–236, Jul. 2012.

[12] N. Elmqvist, A. V. Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. J. Jankun-Kelly, “Fluid Interaction for Information Visualization,” *Information Visualization*, vol. 10, no. 4, pp. 327–340, Oct. 2011.

[13] K. A. Ericsson and H. A. Simon, “Verbal reports as data.,” *Psychological review*, vol. 87, no. 3, p. 215, 1980.

[14] R. J. Gerrig, P. G. Zimbardo, A. J. Campbell, S. R. Cumming, and F. J. Wilkes, *Psychology and life*. Pearson Higher Education AU, 2011.

[15] M. Golemati, C. Halatsis, C. Vassilakis, A. Katifori, and G. Lepouras, “A Context-Based Adaptive Visualization Environment,” in *Tenth International Conference on Information Visualization*, 2006, pp. 62–67.

[16] L. Grammel, M. Tory, and M.-A. Storey, “How Information Visualization Novices Construct Visualizations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 943–952, Nov. 2010.

[17] T. M. Green, W. Ribarsky, and B. Fisher, “Visual analytics for complex concepts using a human cognition model,” in *IEEE Symposium on Visual Analytics Science and Technology*, 2008, pp. 91–98.

[18] S. Greene, E. Tanin, C. Plaisant, B. Shneiderman, L. Olsen, G. Major, and S. Johns, “The end of zero-hit queries: query previews for NASA’s Global Change Master Directory,” *International Journal on Digital Libraries*, vol. 2, no. 2–3, pp. 79–90, Sep. 1999.

[19] H. Guo, S. Gomez, C. Ziemkiewicz, and D. Laidlaw, “A Case Study Using Visualization Interaction Logs and Insight,” *IEEE Transactions on Visualization and Computer Graphics*, vol. PP, no. 99, pp. 1–1, 2015.

[20] L. Harrison, D. Skau, S. Franconeri, A. Lu, and R. Chang, “Influencing Visual Judgment Through Affective Priming,” in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2013, pp. 2949–2958.

[21] J. Heer and M. Agrawala, “Software Design Patterns for Information Visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 853–860, Sep. 2006.

[22] J. Heer, F. van Ham, S. Carpendale, C. Weaver, and P. Isenberg, “Creation and Collaboration: Engaging New Audiences for Information Visualization,” in *Information Visualization*, A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, Eds. Springer Berlin Heidelberg, 2008, pp. 92–133.

[23] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala, “Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1189–1196, Nov. 2008.

[24] J. Heer and G. Robertson, “Animated Transitions in Statistical Data Graphics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1240–1247, Nov. 2007.

[25] D. Kahneman, *Thinking, Fast and Slow*, Reprint edition. New York: Farrar, Straus and Giroux, 2013.
[26] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer, "Enterprise Data Analysis and Visualization: An Interview Study," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2917–2926, Dec. 2012.
[27] D. Kirsh, "Interaction, External Representation and Sense Making," *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, pp. 1103–1108, 2009.
[28] G. Klein, B. Moon, and R. R. Hoffman, "Making Sense of Sensemaking 2: A Macrocognitive Model," *IEEE Intelligent Systems*, vol. 21, no. 5, pp. 88–92, Sep. 2006.
[29] A. Kobsa, "An empirical comparison of three commercial information visualization systems," in *IEEE Symposium on Information Visualization*, 2001, pp. 123–130.
[30] B. chul Kwon, B. Fisher, and J. S. Yi, "Visual analytic roadblocks for novice investigators," in *IEEE Conference on Visual Analytics Science and Technology (VAST)*, 2011, pp. 3–11.
[31] H. Lam, "A Framework of Interaction Costs in Information Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1149–1156, Nov. 2008.
[32] S. Lee, S.-H. Kim, Y.-H. Hung, H. Lam, Y. Kang, and J. S. Yi, "How do People Make Sense of Unfamiliar Visualizations?: A Grounded Model of Novice 's Information Visualization Sensemaking," *IEEE Transactions on Visualization and Computer Graphics*, vol. PP, no. 99, pp. 1–1, 2015.
[33] Z. Liu and J. Heer, "The Effects of Interactive Latency on Exploratory Visual Analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2122–2131, Dec. 2014.
[34] Z. Liu, N. Nersessian, and J. Stasko, "Distributed Cognition as a Theoretical Framework for Information Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1173–1180, 2008.
[35] Z. Liu and J. T. Stasko, "Mental Models, Visual Reasoning and Interaction in Information Visualization: A Top-down Perspective," *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 999–1008, Nov. 2010.
[36] J. Mackinlay, P. Hanrahan, and C. Stolte, "Show Me: Automatic Presentation for Visual Analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1137–1144, Nov. 2007.
[37] J. Maeda, *The Laws of Simplicity*. MIT Press, 2006.
[38] T. Munzner, "A Nested Model for Visualization Design and Validation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 6, pp. 921–928, Nov. 2009.
[39] D. A. Norman, *The Design of Everyday Things*. Basic books, 2002.
[40] H. L. O'Brien and E. G. Toms, "The development and evaluation of a survey to measure user engagement," *Journal of the American Society for Information Science and Technology*, vol. 61, no. 1, pp. 50–69, Jan. 2010.
[41] A. Perer and B. Shneiderman, "Systematic Yet Flexible Discovery: Guiding Domain Experts Through Exploratory Data Analysis," in *Proceedings of the 13th International Conference on Intelligent User Interfaces*, New York, NY, USA, 2008, pp. 109–118.
[42] P. Pirolli and S. Card, "Information foraging," *Psychological Review*, vol. 106, no. 4, pp. 643–675, 1999.
[43] N. H. Riche, B. Lee, and C. Plaisant, "Understanding Interactive Legends: a Comparative Evaluation with Standard Widgets," *Computer Graphics Forum*, vol. 29, no. 3, pp. 1193–1202, Jun. 2010.
[44] J. C. Roberts, "State of the Art: Coordinated Multiple Views in Exploratory Visualization," in *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization, 2007. CMV '07*, 2007, pp. 61–71.
[45] S. F. Roth, J. Kolojejchick, J. Mattis, and J. Goldstein, "Interactive Graphic Design Using Automatic Presentation Knowledge," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1994, pp. 112–117.
[46] P. Saraiya, C. North, and K. Duca, "An insight-based methodology for evaluating bioinformatics visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 4, pp. 443–456, Jul. 2005.
[47] B. Schwartz and K. Kliban, *The paradox of choice: Why more is less*. Ecco New York, 2004.
[48] V. Setlur and M. C. Stone, "A Linguistic Approach to Categorical Color Assignment for Data Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 698–707, Jan. 2016.
[49] B. Shneiderman, "Dynamic queries for visual information seeking," *IEEE Software*, vol. 11, no. 6, pp. 70–77, Nov. 1994.
[50] B. Shneiderman, "The eyes have it: a task by data type taxonomy for information visualizations," in *IEEE Symposium on Visual Languages*, 1996, pp. 336–343.
[51] B. Shneiderman, "Direct Manipulation for Comprehensible, Predictable and Controllable User Interfaces," in *Proceedings of the 2Nd International Conference on Intelligent User Interfaces*, New York, NY, USA, 1997, pp. 33–39.
[52] Y. B. Shrinivasan and J. J. van Wijk, "Supporting the Analytical Reasoning Process in Information Visualization," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2008, pp. 1237–1246.
[53] B. Steichen, C. Conati, and G. Carenini, "Inferring Visualization Task Properties, User Performance, and User Cognitive Abilities from Eye Gaze Data," *ACM Trans. Interact. Intell. Syst.*, vol. 4, no. 2, p. 11:1–11:29, Jul. 2014.
[54] C. Stolte and P. Hanrahan, "Polaris: A System for Query, Analysis and Visualization of Multi-Dimensional Relational Databases," in *IEEE Symposium on Information Vizualization*, Washington, DC, USA, 2000, p. 5–.
[55] J. J. Thomas and K. A. Cook, Eds., *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. Los Alamitos, CA: National Visualization and Analytics Ctr, 2005.
[56] E. Tufte, *The Visual Display of Quantitative Information*, 2nd edition. Cheshire, Conn: Graphics Pr, 1983.
[57] J. Walny, S. Huron, and S. Carpendale, "An Exploratory Study of Data Sketching for Visual Representation," *Computer Graphics Forum*, vol. 34, no. 3, pp. 231–240, Jun. 2015.
[58] J. J. van Wijk, "The value of visualization," in *IEEE Visualization*, 2005, pp. 79–86.
[59] W. Willett, J. Heer, and M. Agrawala, "Scented Widgets: Improving Navigation Cues with Embedded Visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1129–1136, Nov. 2007.
[60] "The Design Ethos of Dieter Rams," *San Francisco Museum of Modern Art*. [Online]. Available: http://www.sfmoma.org/about/press/press_exhibitions/releases/880. [Accessed: 10-Feb-2015].