

Utilizing Provenance as an Attribute for Visual Data Analysis: A Design Probe with ProvenanceLens

Arpit Narechania, Shunan Guo, Eunyee Koh, Alex Endert, Jane Hoffswell

Abstract—Analytic provenance can be visually encoded to help users track their ongoing analysis trajectories, recall past interactions, and inform new analytic directions. Despite its significance, provenance is often hardwired into analytics systems, affording limited user control and opportunities for self-reflection. We thus propose modeling provenance as an attribute that is available to users during analysis. We demonstrate this concept by modeling two provenance attributes that track the *recency* and *frequency* of user interactions with data. We integrate these attributes into a visual data analysis system prototype, ProvenanceLens, wherein users can visualize their interaction recency and frequency by mapping them to encoding channels (e.g., color, size) or applying data transformations (e.g., filter, sort). Using ProvenanceLens as a design probe, we conduct an exploratory study with sixteen users to investigate how these provenance-tracking affordances are utilized for both decision-making and self-reflection. We find that users can accurately and confidently answer questions about their analysis, and we show that mismatches between the user’s mental model and the provenance encodings can be surprising, thereby prompting useful self-reflection. We also report on the user strategies surrounding these affordances, and reflect on their intuitiveness and effectiveness in representing provenance.

Index Terms—analytic provenance, visualization, visual encodings, direct manipulation, visual data analysis, design probe

I. INTRODUCTION

Analytic provenance records the history of analytical actions, showing how data was obtained, transformed, and analyzed. For data visualization, analytic provenance also tracks how users interact with visualizations as a representation of their reasoning process [1]. A frequent use of provenance is to help users recall steps taken during analysis [2]. While effective for forensic purposes, other tools have explored how to show provenance to users during analysis. For example, existing systems [3]–[5] leave visual traces of the user’s interactions to encourage them to pause and reflect on their behavior, potentially influencing subsequent analysis. However, analytic questions such as “*How many data points has the user interacted with so far?*” or “*Which were the first attributes that the user interacted with?*” are often only answerable post-analysis (after analyzing the provenance logs). In the moment, it is impractical for the user to manually count the interaction traces one-by-one or outright remember their interaction history in detail, which makes such iterative reflection during analysis difficult [6]–[8]. Thus, our first research question asks: “*How can we make (such) provenance information available*

Arpit Narechania is with The Hong Kong University of Science and Technology (arpit@ust.hk); this work was done in part when he was with The Georgia Institute of Technology. Alex Endert is with The Georgia Institute of Technology (endert@gatech.edu). Shunan Guo, Eunyee Koh, and Jane Hoffswell are with Adobe Research ({sguo, eunyee, jhoffswell}@adobe.com).

to the user during visual data analysis?” While prior work has already explored modeling [9] and presentation [3] of analytic provenance, this work aims to *standardize* these aspects for consistent and flexible access during analysis.

For modeling provenance, we introduce two **provenance attributes**, *frequency* and *recency*. These attributes track how often and how recently a user interacts with each data attribute and record, assigning a score that is standardized to a common scale from 0.0 (no interaction) to 1.0 (most frequent or most recent interaction, respectively). By modeling the *frequency* of interactions this way, analysts can quickly identify unexplored data, or ones that they keep coming back to for subsequent decision-making. *Recency* on the other hand lets users track when data are interacted with and review how their analytic behavior evolves over time.

For presenting the modeled provenance to users, we propose to dynamically map provenance attributes to one or more visual encoding channels. For example, users can map *frequency* to the color and/or size channels, creating a visualization where data points are colored and/or sized based on their frequency scores, providing a summary of past interactions and influencing subsequent ones. This flexibility contrasts with prior work, where provenance is often *permanently mapped* to specific encodings (e.g., color [3], [10], size [10], opacity [4], or stroke width [4]), preventing users from temporarily disabling them or mapping other information to those channels.

In addition to visual encodings, we also propose sorting and filtering (i.e., transforming data) by provenance attributes to help users strategize about subsequent interactions, such as organizing the data, e.g., “rank data attributes by their frequency” (sort) or reducing the search space, e.g., “show top five most recently interacted points” (filter). Collectively, we standardize provenance visualization for data attributes and records while also affording more user flexibility to encode, filter, and sort provenance, or combinations thereof.

Such modeling and presentation of provenance information motivates our second, central research question, “*How do people use the provenance attributes during visual data analysis?*” For example, do users prefer mapping provenance to color, size, both color and size, or something different, like an axis? If not visual encodings, do users prefer interacting with provenance via data transformations such as sorting or filtering? Lastly, do these user preferences change for specific tasks, e.g., while reviewing someone else’s analysis history versus doing one’s own analysis?

To study the utility and usage characteristics of provenance attributes during visual analysis, we developed a system prototype, **ProvenanceLens**, that allows users to map provenance

to visual encodings and data transformations; users can now interact with provenance in the same way as regular attributes.

Using ProvenanceLens as a design probe, we conducted an **exploratory study** with sixteen users to investigate how they utilized provenance attributes for a decision-making task. Specifically, we tasked users to (1) *review* and *answer* questions about another user’s provenance that is preloaded into the system and then (2) *track* and *answer* questions about their own provenance. Our users also self-reported their confidence and surprise levels while answering these questions.

We found that provenance attributes in ProvenanceLens can help users accurately and confidently answer questions about another user’s as well as their own analysis. Users were often surprised and uncovered unexpected insights and analytical behaviors that they may have misremembered, facilitating self-reflection. In addition to user performance, we report on user preferences and strategies for utilizing available affordances, assessing their intuitiveness and effectiveness to track and recall analytic provenance. For instance, we found that *x*, *y*, and *fill* were the most preferred encodings, whereas, *shape* and *stroke* were either rarely used or not used at all. We conclude by discussing design implications for future provenance-based systems and studies. Our primary contributions include:

- 1) Modeling and presentation of the *frequency* and *recency* of **analytic provenance as attributes** (i.e., “**provenance attributes**”) during visual data analysis.
- 2) Findings from an exploratory study with sixteen users on **how they use provenance attributes in a visual analysis system** to accurately and confidently answer questions about their analysis, sometimes causing surprise.

II. RELATED WORK

A. Analytic Provenance: Tracking, Modeling, and Visualizing

Our memory has a finite capacity to track and remember our past interactions with data, which affects exploration [6], [7]. Analyzing prior data interactions in a visualization is a form of analytic provenance [1], [2] that is often used to infer one’s analysis process. Studying such analytic provenance has a well-documented history in visualization research. Heer et al. have summarized the design space for displaying interaction histories [11]. Xu et al. [12]’s survey notes that provenance is broadly used to support sensemaking [13], [14] and decision-making [15] with data, evaluate the usefulness of visualization systems [16], [17], design adaptive systems [18], [19], improve the performance of machine learning models [20], replay or replicate analysis sessions [21], [22], and automatically generate summary reports of an analysis session [23], [24]. Provenance also increases awareness of analytic behavior [3], mitigates human biases [25], increases unique data discoveries [4], [5], and affects confidence levels [26] during analysis.

Today, provenance tracking occurs in various contexts including analysis tools [1], code editors [27], computational notebooks [28], [29], workflow modeling systems [21], [30], collaborative environments [31]–[33], websites [34], [35], and video games [36], [37]. There are several logging frameworks that help capture this provenance [38]–[41]. Relevant to this work on visual data analysis, provenance has been

modeled from user interactions (via mouse [42] or gaze [43] tracking) with UI controls [10], [44]–[46], visualizations as a whole [4], [8], [47], [48] or with specific elements and regions within them (e.g., marks) [1], [2], [4], including attributes and records [3], [25]. Several metrics have been proposed that quantify analytic focus [9], exploration behavior [49] and subsequent interaction predictions [50], differences between a user’s subset selections and the data [51], and differences between a user’s interactions and a baseline [52]. Our work extends previous efforts to model provenance from mouse-based interactions with data by proposing a more standardized approach (i.e., the concept of “provenance attributes”) that facilitates flexible, real-time user access to this provenance during analysis.

Lastly, provenance visualizations often involve encoding provenance on or near the object of interaction, e.g., highlighting previously visited or interacted data attributes and records [3], visualizations as a whole [4], [32], [33] or specific regions within them [53], regions and hyperlinks on a webpage [35], [54], options and ranges in UI controls [5], [10], lines of code in a code editor [27], or a document’s authorship and readership history [55], [56], among others. Other provenance visualizations are separately visualized in an external view or application, e.g., as a graph [21], [41], [57]–[59]. Relevant to our work, Lumos [3] tracks the frequency of user interactions with data and leaves behind visual traces (using color) to increase user awareness of analytic behaviors. Our work extends Lumos [3] to track not only the frequency but also the recency of interactions with data, and to present provenance via not only color but also a variety of visual encodings (e.g., size, shape, tooltip) and data transformations.

B. Direct Manipulation

Direct manipulation is an interaction technique wherein objects of interest in the UI are visible and users can manipulate them through physical, reversible, and incremental actions and receive immediate feedback [60]. This technique has been effectively used to enhance user interaction and comprehension during visual data analysis [61]–[63]. For instance, Tableau [64] allows users to perform drag-and-drop operations to specify data visualizations and apply data transformations. Wrangler [65] utilizes direct manipulation to help users author expressive transformations while simplifying specification and minimizing manual repetition. Relevant to our work, DataPilot [66] helps users to navigate large, unfamiliar datasets by enabling them to sort and filter data attributes and records based on their quality and usage characteristics [67]. We similarly enable users to sort and filter data but based on provenance (i.e., recency, frequency of their interactions) during analysis.

III. UTILIZING PROVENANCE AS AN ATTRIBUTE DURING VISUAL DATA ANALYSIS

In this section, we describe how we (A) track user interactions with data attributes and records to (B) model provenance attributes and subsequently (C) visualize them during analysis.

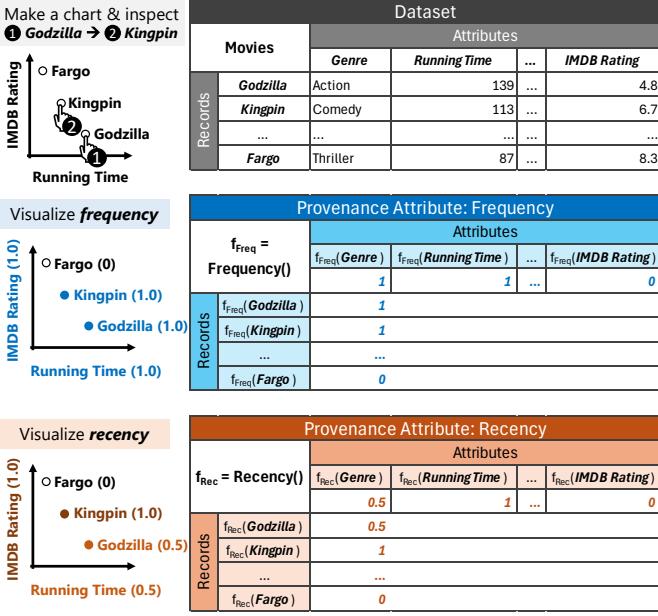


Fig. 1. Illustration of two provenance attributes, **frequency** and **recency**, modeled for each attribute and record of a dataset about movies, and scored on a scale from 0 (low) to 1 (high). Consider a user creates a scatterplot visualization of (1) **Running Time** \times (2) **IMDB Rating** and then clicks two datapoints one after another (1) **Godzilla** \rightarrow (2) **Kingpin**, indicating interactions with two attributes and two records. Regarding data attributes, **IMDB Rating** and **Running Time** both receive a **frequency** score of 1.0 (each interacted once, hence maximum score), while other attributes score 0.0; for **recency**, **IMDB Rating** (most recently interacted) scores 1.0 and **Running Time** scores 0.5 (user mapped *x* axis before *y* axis), while other attributes score 0.0. Likewise, regarding data records, **Godzilla** and **Kingpin** both score 1.0 on **frequency**; but for **recency**, **Kingpin** (most recently interacted) scores 1.0 and **Godzilla** scores 0.5, while other records score 0.0. These scores are derived by evenly spacing the interactions between 0 and 1, based on their count and order of occurrence in the interaction history.

A. Tracking Provenance: Which User Interactions to Log?

In a typical visual data analysis system, users analyze their dataset in various ways: they can inspect an attribute's summary statistics (e.g., via distribution plots), examine individual data records (e.g., from a data table), apply data transformations (e.g., filter and sort), and create visualizations (e.g., by mapping attributes to visual encodings). To achieve our overarching goal of tracking, modeling, and visualizing provenance during analysis, we track a subset of relevant user interactions and map them to an individual data attribute or record. In this work, we track interactions with a data *attribute* when a user inspects its summary profile, maps it to a visual encoding, or uses it to filter and sort data records; we track interactions with a data *record* when a user hovers on a visualization mark (e.g., a point in a scatterplot) or a row in the data table.

B. Modeling Provenance Attributes: Frequency, Recency

After tracking user interactions, accurately modeling provenance attributes is crucial for understanding analytic behaviors. To model provenance in our work, we generally follow the methodology employed by Lumos [3]. For unit visualizations, we map one interaction with an attribute or record as +1 unit of interaction. For aggregate visualizations that show a single value computed from multiple data records (e.g., a bar in a bar chart with an aggregation function such as *mean* or *sum*),

we map one interaction with an aggregated entity as $+1/N$ units of interaction for each of the N data records that form the hovered entity. For example, consider a bar chart showing average **IMDB Rating** for different movie *Genres*; if the user hovers on the bar *Genre*=“Action” that represents five action movies (whose mean is encoded as the bar’s height), we log each action movie as having +0.2 units of interaction.

We also log the timestamp (as milliseconds since epoch) of each interaction. For unit visualizations, we simply map the interaction timestamp to the corresponding attribute or record. For aggregate visualizations, we log the same interaction timestamp for all data records that form the aggregated entity (i.e., many records will have the same interaction timestamp).

From the interaction units and timestamps, we compute two metrics that we refer to as **provenance attributes**: **frequency** and **recency**. We chose to focus on these metrics as they are both relevant to provenance tracking and have been commonly used in visualization research and practice [3], [9], [10], [25], [41]. Figure 1 illustrates sample **frequency** and **recency** computations, which are described in the following paragraphs.

Frequency. This provenance attribute computes a frequency score f_x normalized from zero to one, for each data attribute:

$$f_x = \frac{n_x}{\max_{i=1}^N n_i}$$

where n_x is the total number of interactions with a data attribute x and $\max_{i=1}^N n_i$ is the maximum number of interactions from among all N data attributes. A score of zero implies no interactions (or zero focus) and a score of one implies the most number of interactions (or maximum focus). Like data attributes, we also compute f_x for data records.

Recency. This provenance attribute computes a recency score r_x normalized from zero to one, for each data attribute:

$$r_x = \frac{\text{rank}_N(\max(t_x))}{\sum_{i=1}^N n_i}$$

where $\max(t_x)$ is the timestamp of the most recent interaction with a data attribute x , $\text{rank}_N(\max(t_x))$ is its serial order in the overall sequence of interactions across all N data attributes (i.e., the entire analysis history), and $\sum_{i=1}^N n_i$ is the total number of interactions across all N data attributes. A score of zero implies no interactions (or zero focus) and a score of one corresponds to the most recent interaction (or focus). Like data attributes, we also compute r_x for data records.

C. Visualizing & Interacting with Provenance during Analysis

Our main goals were to enable users to access the provenance of specific attributes or records and also obtain a visual provenance overview of the entire dataset during analysis. In response, we designed small glyphs called “**provenance attribute glyphs**”, which comprise a mark type (e.g., *point* ○) and one or more visual encodings (e.g., *fill* ●) that encode the provenance attribute values (e.g., ●●● where darker glyphs imply higher values). These glyphs can represent data records within visualizations (e.g., points in a scatterplot), can be displayed alongside data attributes (e.g., in the attribute panel), and can integrate well with sorting and filtering operations.

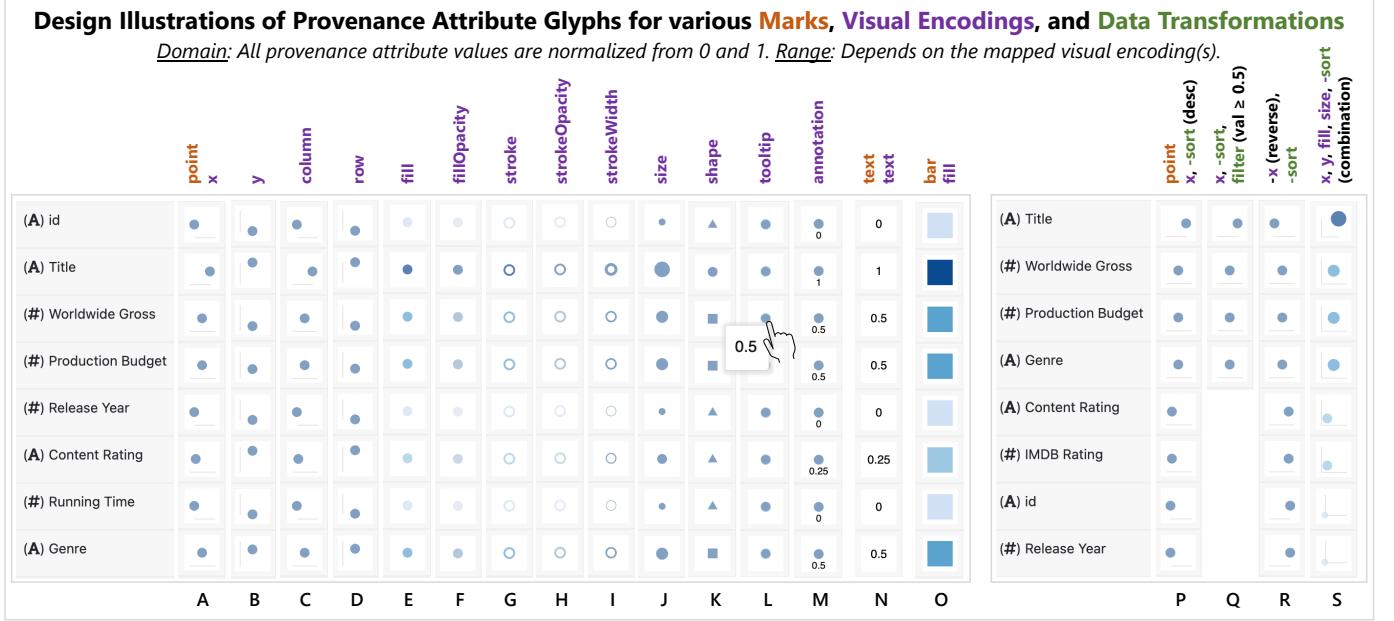


Fig. 2. Design illustrations of **provenance attribute glyphs** (A–S) for data attributes (or records) across different **marks** (point, text, bar), **visual encodings** (x, y, column, row, fill, fillOpacity, stroke, strokeOpacity, strokeWidth, size, shape, tooltip, annotation, text), and **data transformations** (sort, filter), including alternate configurations (e.g., $-x$ where the range is reversed or descending sort order) and combinations (e.g., $x + y + fill + size + sort$). For instance, for **mark=bar** and **encoding=fill** (O): “Title” **■** has the largest value (darkest bar) followed by “Worldwide Gross” **■**, “Production Budget” **■**, and “Genre” **■**; “id” **■**, “Release Year” **■**, and “Running Time” **■** have the smallest values (lightest bars). Notice the change to the attribute sort order for the right side of the figure (P–S), compared to the unsorted attributes on the left (A–O).

These glyphs can reveal interesting provenance patterns, e.g., if there are more dark than light glyphs (or vice versa), then the user has interacted disproportionately. Figure 2 illustrates the supported designs of provenance attribute glyphs for data attributes, described in detail below.

1) *Marks and Visual Encodings:* We cover three mark types: *point* \circ , *bar* \blacksquare , and *text* (0.5). Other mark types such as *line* and *area* require at least two values (a start and an end), which make them unsuitable for encoding, and hence are not considered. Each mark type encodes a single value across one or more visual encodings, described next.

We cover thirteen encodings [68]: *x*, *y*, *column*, *row*, *fill*, *fillOpacity*, *stroke*, *strokeOpacity*, *strokeWidth*, *size*, *shape*, *tooltip*, *text*, and *annotation*; *annotation* is a special encoding that adds an extra text mark displaying the encoded value next to the visualization mark, unlike the *text* encoding, that only displays the text (as the visual mark itself).

For instance, Figure 2O (**mark=bar**, **encoding=fill**) shows that “Title” has the largest value (darkest bar \blacksquare), whereas “id”, “Release Year”, and “Running Time” have the smallest values (lightest bar \square). Similarly, Figure 2I (**mark=point** and **encoding=strokeWidth**) shows the glyph \circ for the largest value (thick stroke) and \circ for the smallest value (thin stroke).

2) *Data Transformations:* In addition to visualizing provenance attributes (marks and encodings), we also enable users to transform (filter and sort) their data by the provenance attributes. This functionality was inspired by DataPilot [66], which similarly lets users sort and filter their data attributes and records based on their quality and usage characteristics.

Sort displays the data attributes or records in order of the encoded provenance attributes. For example, Figure 2P

shows an active descending **-sort** by *frequency*, visualized as **mark=point** and **encoding=x**, illustrating that “Title” had the most interactions followed by “Worldwide Gross”, “Production Budget”, and so on. Notice the change to the attribute sort order for the right side of Figure 2 (P–S), compared to the unsorted attributes on the left (A–O).

Filter displays a subset of data attributes or records based on the criteria provided by the user (e.g., the encoded entity values). For example, Figure 2Q shows an active **filter** for *frequency* greater than or equal to 0.5, emphasized by **mark=point**, **encoding=x**, and a descending **-sort** (also by *frequency*), resulting in four attributes that match the filter (“Title”, “Worldwide Gross”, “Production Budget”, “Genre”).

3) *Configurations and Combinations:* We enable various configurations and combinations of marks, visual encodings, and data transformations, affording user agency and control, and promoting accessibility during analysis.

Configurations. Users can configure the range of the encoding scale (e.g., for *size*, high values map to bigger glyphs or vice versa) or sort directions (i.e., ascending or descending). For example, all data attributes to the right side of Figure 2 (P–S) are **-sorted** by *frequency* in the descending order (which means attributes at the top have been used more often in the interface). However, **P** maps the glyphs to **x** whereas **R** maps them to **-x** (reverse). Such configurability can support different user preferences and use cases, e.g., visited points could become smaller to nudge the user to interact with other points (increase coverage); alternatively, they could become bigger to help the user quickly spot them.

Combinations. We do not limit the user to one configuration at a time. Users can simultaneously utilize one or

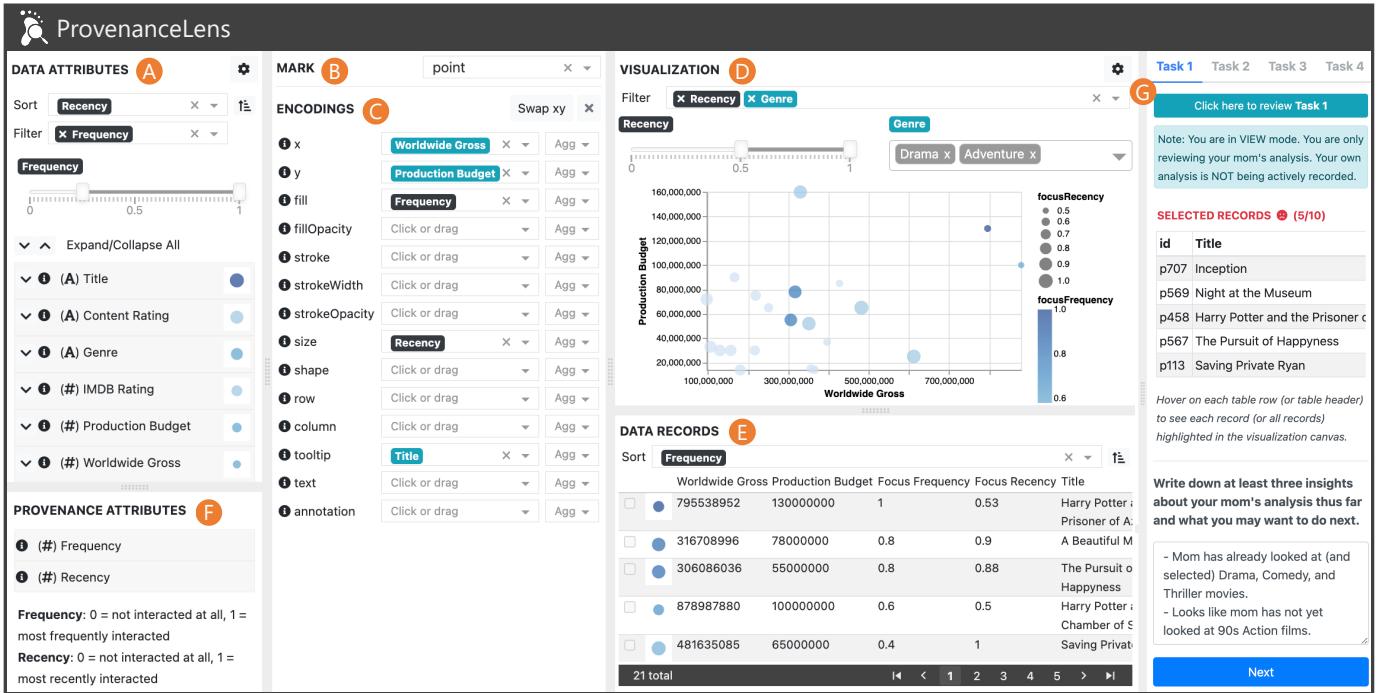


Fig. 3. The ProvenanceLens user interface consisting of seven views: (A) the **Data Attributes** view shows the attributes and enables transformation (e.g., sort, filter); (B) the **Marks** and (C) **Encodings** views specify the visualization; (D) the **Visualization** view renders the specified visualization and supports filtering of data records; (E) the **Data Records** view supports review and transformation (sort) of the data records shown in the visualization; (F) the **Provenance Attributes** view lists the *recency* and *frequency* attributes; and (G) the **Tasks** view shows the task instructions and questions, and tracks the user's progress.

more visual encoding assignments, as well as filter and sort criteria to realize a wide variety of glyph combinations. For example, Figure 2S shows a combination of *x*, *y*, *fill*, *size*, and *sort*. Combinations can help reinforce certain interaction patterns and also help tools ship with smart defaults to serve a wider audience, e.g., using both *fill* and *size* can support both colorblind and non-colorblind users [69], similar to how figures in articles are often colored but also hatched.

IV. EVALUATION: EXPLORATORY USER STUDY

To study the utility and usage characteristics for provenance attributes during visual data analysis, we designed a decision-making task that involves reviewing and answering questions about both another user's and one's own analytic provenance. In this section, we first describe this prototype, ProvenanceLens, as well as our overall study design(s).

A. ProvenanceLens: System Prototype for the User Study

With no similar existing system, we developed one wherein users can utilize provenance attributes to perform visual data analysis (e.g., inspect a dataset, create visualizations, and apply transformations) and answer study-specific questions about their provenance. Figure 3 shows the UI with seven views:

A Data Attributes. In this view, users upload and configure the dataset (⚙️) and see the underlying attributes (or features or columns). Hovering on the information icon ⓘ shows the attribute's definition in a tooltip. Clicking on the expand icon ⓘ opens a detailed view with a distribution plot of the attribute's values: an area curve for numerical attributes and a column chart for categorical attributes, both of which show

percentage counts corresponding to the attribute quantiles and categories, respectively. Users can sort and filter the attributes using the *recency* and/or *frequency* provenance attributes.

B Mark. This view includes a dropdown to configure the mark type for the visualization. Users can select one of *point*, *bar*, *line*, *area*, or *text* to begin a visualization specification.

C Encodings. This view shows the visual encoding channels. Users select or drag one or more attributes (data and/or provenance) to one of *x*, *y*, *fill*, *fillOpacity*, *stroke*, *strokeOpacity*, *strokeWidth*, *shape*, *row*, *column*, *tooltip*, and/or *text* to complete a visualization specification. An additional encoding, *annotation*, adds a new annotation displaying the value of the encoded entity next to the selected mark.

D Visualization. This view renders an interactive visualization based on the selected mark type and activated visual encodings in the “Encodings” view. It also includes a “Filter” drop-zone to filter out data points by attributes (data and/or provenance). A numerical attribute displays a range slider and a categorical attribute displays a multi-select dropdown.

E Data Records. This view shows the data bound in the visualization as a paginated data table. If the user hovers on a data point in a unit visualization (e.g., a scatterplot), this table filters to only show that data record whereas if the user hovers on an entity in an aggregate visualization (e.g., a bar showing the *mean* value), this table filters to show all data records belonging to the hovered entity (e.g., the bar).

F Provenance Attributes. This view shows the two provenance attributes: *frequency* and *recency*. Like data attributes, users can select or drag these provenance attributes and drop them to the sort and/or filter drop zones in the “Attribute

View”, the encoding channels in the “Encodings” view, the filter drop zone in the “Visualization” view, or the sort drop zone in the “Data Records” view.

G Tasks. This view has six tabs, one for each task (T1–T6) for study participants to access the current task instructions, track their progress, and answer questions via integrated forms.

ProvenanceLens can be programmatically configured into three modes: (1) *edit*-only, wherein provenance is tracked and presented back to the user in real-time, (2) *view*-only, wherein existing provenance is imported into ProvenanceLens without real-time tracking, and (3) *hybrid*, wherein provenance is imported into the system, and real-time tracking is also enabled.

B. Example Usage Scenarios for ProvenanceLens

We present two usage scenarios on how ProvenanceLens can enhance visual data analysis for real-time provenance tracking as well as post-analysis review of a user’s provenance.

Real-time Provenance Tracking. Assume Mark works for a movie production company and must determine what kinds of movies to make next. They upload the dataset of movies in ProvenanceLens (configured in its *edit*-only mode, i.e. real-time provenance tracking) and begin exploring (by specifying different visualizations, applying relevant filters, and hovering on certain movies). After taking a short break, they wish to revisit their most recently hovered movie, so they create a visualization with a *point* mark type and map the *recency* provenance attribute to the *x* encoding. They immediately find their desired movie (with the highest recency score) at the horizontal axis’ rightmost-end. Had they mapped *recency* to another encoding such as *fill* or *size*, it may have taken them some time to accurately determine the darkest or biggest point.

Upon further inspection of their focus on different movies, Mark noticed they have interacted with some movies multiple times and not considered a lot of other movies (i.e., they exhibited less exploration coverage). Mark wants to change this analytic behavior, and thus starts visualizing traces of their interactions in real-time. Because they are colorblind, they avoid the *fill*, *stroke*, *fillOpacity*, and *strokeOpacity* encodings and choose *size* to encode the *recency* provenance attribute. They can now track visited points using their bigger size and continue exploring. However, because visited points get bigger, they are getting drawn to the same points even more; thus, they reverse the *size* range to make the visited points smaller instead. Happy with this configuration, they continue exploring and eventually submit their analysis report to their manager.

Post-Analysis Provenance Review/Audit. In addition to real-time provenance tracking, ProvenanceLens can also *import* an existing log of a user’s analytic provenance to facilitate collaboration (e.g., continuing a colleague’s analysis), auditing (e.g., inspecting a colleague’s analysis), or post-hoc analysis (e.g., reviewing user study logs [30], [70]).

Assume Anya is Mark’s manager and is reviewing their previous analysis. They express surprise at Mark’s recommendation to make a *Drama* movie next. Wanting to review Mark’s analysis, Anya imports Mark’s analytic provenance, that was exported from ProvenanceLens, back into ProvenanceLens.

They configure ProvenanceLens in its *view*-only mode, i.e., no real-time provenance tracking, to review and interact with Mark’s analysis. They make a bar chart with “Genre” on *x* and the sum of *frequency* on *y*. This visualization shows Mark’s total focus across movie genres. Anya notices that certain bars corresponding to *Action* and *Adventure* movies are really short (i.e., of low *frequency*). They ask Mark to review those genres before making a final recommendation. In this way, Anya was not only able to review Mark’s report but also their analysis process, thereby supporting enhanced decision-making.

C. Pilot Studies and Evaluation Considerations

Before finalizing our study design, we explored two alternate designs and also conducted subsequent pilot studies.

Pilot Study 1: Decision-Making. We recruited four Ph.D. students (three years into the program) as our pilot users. We tasked users to explore a dataset of movies and select (1) ten movies (records) satisfying certain criteria, (2) four movie characteristics (attributes) that were important to their analysis, and (3) answer a series of questions about their analysis.

We observed users did not utilize the provenance attributes to track their analysis process and only used them during the subsequent question-answering. We also noticed that answering questions immediately after analysis may not be hard for certain users because the analysis may be quite fresh in their memory. In addition, the idea to make users select important attributes was straightforward and seemed redundant because the task criteria already hinted what attributes to consider.

Pilot Study 2: Analysis Review and Decision-Making. We recruited two other Ph.D. students (four years into their program) as our pilot users. In this revised study design, we introduced an initial *review* task to first make the user explore another user’s analytic provenance, write three insights, and then answer questions about the prior analysis process. The original task to select movies followed this task. Our goal was to make the user actively utilize the provenance attributes during their own analysis, and we believed making them answer questions about another user’s analysis first would remove the fresh-in-memory aspect from the equation. We also discarded the selection of four movie characteristics. Overall, we noticed this study design had a desired positive effect and decided to use it for our eventual user study, described next.

D. Exploratory User Study Design

Participants: We recruited 16 participants from a public university in the U.S. who were pursuing a bachelors (1), masters (11), or doctoral (4) degree in computing or related fields (15) and economics (1). These participants were either enrolled in or alumni of at least one data visualization class and self-reported their visualization literacy to be at least 3 on a scale from 1 (novice) to 5 (expert). Demographically, they were in the 18-24 (7) or 25-34 (9) age groups (in years) and of *female* (5), *male* (10), or *preferred not to say* (1) genders.

Study Session: Each study session lasted between 75 and 90 minutes. We compensated each participant with a \$15 gift card for their time. We conducted the study remotely over Zoom;

the experimenter provided participants access to the study environment by sharing their (the experimenter's) computer screen and granting input control to the participant. After providing consent, participants saw a five-minute video tutorial that demonstrated the features of ProvenanceLens. Participants then performed a practice task on a *dataset of cars* to get acquainted with the UI before starting the actual task.

The actual task was on a *dataset of movies* and lasted about 60 minutes. Participants were not required to think aloud during the task to simulate a realistic work setting (although some participants felt comfortable doing so). During the task, participants' interactions with the system were logged. The study ended with participants completing a feedback questionnaire and a background questionnaire. Each study session was screen- and audio-recorded for subsequent analysis.

Task and Dataset: We designed the following visual data exploration and decision-making task about a movies dataset:

Imagine your family is planning a month-long vacation to Europe. Going with you are your siblings, parents, grandparents, your uncle and aunt, and their boy (your cousin). Your mom began selecting some movies to pick and watch from for the occasional movie nights. **Her target is to select ten movies to carry to the vacation.**

Because she wanted to ensure a delightful and well-rounded movie night experience for your entire family, **she took the below suggestions and preferences from some of your family members into consideration.**

- 1) Dad: *"I like thrillers and comedies."*
- 2) Uncle: *"Let's watch a heartfelt drama."*
- 3) Cousin: *"I want to re-watch a Harry Potter movie."*
- 4) Grandpa: *"I would love to watch a 90s' action film."*
- 5) Aunt: *"I want to watch a hidden gem: a highly-rated movie that didn't do well commercially."*

After selecting five movies, your mother got pulled to do another task, and assigned you to complete it.

Thus, you will complete the following six tasks (T1-T6).

— Working with your mom's analytic provenance —

T1 Review: Review the five movies your mom selected and many others she explored; write three insights.

T2 Recall: Answer five objective questions about mom's analysis.

T3 Visualize: Create visualizations to answer subjective questions about mom's analysis.

— Working with your own analytic provenance —

T4 Analyze: Select the remaining five movies.

T5 Recall: Answer five objective questions about your analysis.

T6 Visualize: Create visualizations to answer subjective questions about your analysis.

To help you review and track the movies (records) and their characteristics (attributes) one has *interacted* with how many times and when, our system provides two special provenance attributes: **Frequency** and **Recency**. You can use these in the same way you would use the dataset attributes, i.e., create visualizations by mapping them to visual encodings, filter by them, or sort by them.

Note that "*interacted*" refers to a user's interactions in the interface such as hovering on a record to get additional details, mapping dataset attributes and/or provenance attributes to visual encodings (e.g., 'Genre' to *x*), applying a filter (e.g., 'Title'='Titanic') or sort.

V. RESULTS

In this section, we present general and task-specific findings from the user study along with participant performance on

the six tasks (T1-T6) and discuss them in the context of the qualitative feedback from our participants (*P_{1,...,16}*). Detailed participant responses and performance assessments for tasks T1-T6, post-study feedback questionnaire and responses, and a demo video of the UI are available in supplemental material.

A. General Feedback

Participants scored ProvenanceLens 80.94 out of 100 on the system usability scale (SUS [71]), finding it very useful. *P₁₅* acknowledged it is impossible to remember everything during analysis, calling for "*Tools like Tableau [to] really overlay communities' or experts' interaction stats such as frequency and recency in their system. It is one of those features that only adds value and is not necessarily a hindrance.*" *P₁₄* was impressed that the two provenance attributes could help answer a variety of questions, enabling users to revisit, review, and maybe even recreate someone's history of interactions while doing analysis. *P₁₁* liked being able to detect subconscious interaction patterns during analysis. *P₁₆* suggested a potential use case, "*If you are a manager and if you need to review the decision process of your employees, then it is very useful.*"

B. Task Performance

Table I shows the task-specific breakdown of participants' accuracy and self-reported confidence and surprise, along with the number of interactions and time spent during the study.

Accuracy. We computed accuracy by comparing participant's responses to each question with corresponding answers computed from their actual interaction behavior (ground truth). To ensure consistency and reliability, we automated this process using a standardized Python script, as the questions were fixed across all participants. Participants performed exceedingly well overall, achieving high accuracies during the **recall** tasks (T2, T5) and the **visualize** tasks (T3, T6). The mean accuracies were as follows: T2 ($\mu = 98.75\%$), T3 ($\mu = 100\%$), T5 ($\mu = 98.75\%$), T6 ($\mu = 100\%$). We did not compute accuracy for the **review** (T1) and **analyze** (T4) tasks as they were exploratory in nature.

Only two participants (*P_{13,14}*) incorrectly answered one question each. For instance, for T2.Q2 ("Which was the last (i.e. most recent) movie YOUR MOM interacted with?"), *P₁₄* selected the 'point' mark type and assigned "Genre" to *x*, *recency* on *size*, and "Title" on *tooltip*. Then, they applied a *recency* filter of [0.63, 1]. This configuration resulted in four marks stacked on top of each other (unknown to the participant who thought there was only one). The user hovered on it, read "*The Curious Case of Benjamin Button*" in the tooltip, and selected it as the answer. However, this answer was incorrect because "*Saving Private Ryan*", under this movie, had the higher recency score of 1.0. This error may have been avoided if instead of (or along with) *size*, the user assigned *recency* to *x* or *y*, which would have spaced out the points.

Confidence & Success. On a scale from 1 (low) to 7 (high), participants self-reported very high *confidence* in answering questions during tasks T2 ($\mu=6.76$, $M=7$), T3 ($\mu=6.48$, $M=7$), T5 ($\mu=6.70$, $M=7$), and T6 ($\mu=6.48$, $M=7$). In addition, while

| Task | Que | Description | $\mu\text{Int.}$ | μTime | $\mu\text{Acc.}$ | $\mu\text{Conf.}$ | $\mu\text{Sur.}$ |
|----------------|-----|--|------------------|------------------|------------------|-------------------|------------------|
| | | | (mins) | (%age) | (/7) | (/7) | |
| T1 (Review) | Q1 | Review YOUR MOM's analysis and write three insights. | 53.43 | 7.32 | - | - | - |
| T2 (Recall) | Q1 | Select the movie characteristic(s) that YOUR MOM interacted with the MOST. | 12.69 | 1.35 | 100 | 6.81 | - |
| | Q2 | Which was the last (i.e. most recent) movie YOUR MOM interacted with? | 15.88 | 2.32 | 93.75 | 6.81 | - |
| | Q3 | Did YOUR MOM ever interact with the movie 'Titanic'? | 12.69 | 1.89 | 100 | 6.75 | - |
| | Q4 | Did YOUR MOM interact with at least one movie from all 'Content Rating's? | 13.38 | 1.94 | 100 | 6.69 | - |
| | Q5 | Did YOUR MOM interact with all attributes at least once? | 17.88 | 1.49 | 100 | 6.75 | - |
| T3 (Visualize) | Q1 | What was the distribution of YOUR MOM's focus across different movie 'Genre's? | 12.13 | 2.08 | 100 | 6.75 | - |
| | Q2 | How did YOUR MOM's focus on 'Drama' Movies evolve over time? | 35.50 | 4.22 | 100 | 6.06 | - |
| | Q3 | Which were YOUR MOM's most FREQUENTLY interacted movies? Try to show five. | 24.69 | 2.69 | 100 | 6.63 | - |
| T4 (Analyze) | Q1 | Select the five remaining movies. | 89.75 | 7.24 | - | - | - |
| T5 (Recall) | Q1 | Select three movie characteristic(s) YOU most RECENTLY interact with. | 13.75 | 1.90 | 100 | 6.88 | 1.88 |
| | Q2 | Which was the first (i.e. earliest) movie YOU interacted with? | 22.13 | 3.01 | 93.75 | 6.25 | 3.56 |
| | Q3 | Did YOU interact with the movie 'Pearl Harbor'? | 8.63 | 1.44 | 100 | 6.81 | 1.63 |
| | Q4 | Which 'Content Rating' category did YOU interact with the LEAST or NONE AT ALL? | 11.25 | 1.98 | 100 | 6.69 | 4.19 |
| | Q5 | Did YOU interact with at least half of the movie characteristics available? | 12.38 | 1.22 | 100 | 6.88 | 1.94 |
| T6 (Visualize) | Q1 | How similar were YOUR interaction patterns for 'Comedy' and 'Thriller' movies? | 17.81 | 2.67 | 100 | 6.38 | 2.56 |
| | Q2 | Which were YOUR most RECENTLY interacted movies? Try to show THREE movies. | 19.00 | 2.42 | 100 | 6.63 | 2.69 |
| | Q3 | Given an opportunity, which (kinds of) movies would YOU like to go back and interact with? | 18.88 | 2.74 | 100 | 6.44 | 2.38 |

TABLE I

TASKS AND SUMMARY PERFORMANCE STATISTICS FOR SIXTEEN PARTICIPANTS: **Task and Question Index**, **Task Description**, AVERAGE NUMBER OF INTERACTIONS ($\mu\text{Int.}$) AND TIME SPENT IN MINUTES (μTime), AND WHEREVER APPLICABLE, AVERAGE ACCURACY ($\mu\text{Acc.}$), AVERAGE CONFIDENCE ($\mu\text{Conf.}$), AND AVERAGE SURPRISE ($\mu\text{Sur.}$) ON A SCALE FROM ONE (LOW) TO SEVEN (HIGH). T1 AND T4 WERE EXPLORATORY IN NATURE, HENCE WE DID NOT COMPUTE ACCURACIES OR ASK PARTICIPANTS TO SELF-REPORT CONFIDENCE AND SURPRISE SCORES. SIMILARLY, T2 AND T3 WERE FOCUSED ON ANSWERING QUESTIONS ABOUT MOM'S ANALYSIS, HENCE WE DID NOT ASK PARTICIPANTS TO SELF-REPORT THEIR SURPRISE SCORES. OVERALL, PARTICIPANTS PERFORMED EXCEEDINGLY WELL ON ALL THE TASKS, ACHIEVING HIGH ACCURACIES WITH HIGH CONFIDENCE WITH SOME MOMENTS OF SURPRISE.

answering questions based on their own analysis, participants self-reported varying *surprise* during tasks T5 ($\mu=2.64$, $M=2$) and T6 ($\mu=2.54$, $M=2$). Note that T1 and T4 were exploratory in nature, hence we did not ask participants to self-report confidence and surprise scores. Similarly, T2 and T3 were focused on answering questions about mom's analysis, hence we did not ask participants to self-report their surprise scores.

Fidelity. On a scale from 1 (low) to 5 (high), participants self-reported that overall, all six tasks (T1-T6) caused low *physical demand* ($\mu=1.44$, $M=1$), *temporal demand* ($\mu=2.44$, $M=2$), and *frustration* ($\mu=1.5$, $M=1$), average *mental demand* ($\mu=3.31$, $M=3$) and *effort* ($\mu=2.94$, $M=3$), but resulted in high *performance* ($\mu=4.38$, $M=4$). On average, participants spent around 56 minutes completing the study and performed 409 interactions. We did not seek feedback after each task to avoid demotivating the user's analysis with frequent interruptions.

Summary. The high *accuracy*, *confidence*, and *performance* along with average *effort* and *mental demand* suggest that participants were able to effectively use provenance attributes to answer the questions. The high variance in surprise scores suggests that participants were sometimes able to recall their analysis from memory (less surprise) or had an incorrect mental model or recall of their interactions (more surprise).

C. Task-specific Feedback

Reviewing Mom's Previous Analysis (Review: Task T1). Due to the nature of this task, all participants leveraged the provenance attributes and found them to be useful. Participants felt that the provenance attributes could provide "*a record of [mom's] analysis*" (P₁₂), which made it easier to "*track her way of thinking through a quantitative behavior analysis*" (P₃).

For P₆, the provenance attributes were only somewhat useful because they were able to understand mom's analysis simply by inspecting her final selections. On the other hand, P₄ argued that the provenance attributes "*gave context for [mom's movie selections]*," which they would have to otherwise guess (P₁₂).

Answering Questions about Both Mom's and One's Own Analysis (Recall, Visualize: Tasks T2, T3, T5, T6). Like T1, due to the nature of the questions, all participants found the provenance attributes to be useful. P₁ noted, "*It was very useful to use provenance attributes as filters and encodings. I would not have been able to answer questions effectively without the visualization ability.*" P₁₄ noted, "*While looking at and analyzing [mom's] focus, it was hard to keep all the insights in mind, so having it visualized helped understand and hence remember it better.*" P₉ found the provenance attributes to be more helpful to answer questions about visualizing focus (T3, T6) than searching for insights (T2, T5), noting that "*I was able to concentrate on specific facets of [mom's] focus at a time [via the visualizations], which felt more organized.*"

Selecting Remaining Movies (Analyze: Task T4). Ten participants (P_{1,2,3,5,6,7,10,11,12,16}) either did not use the provenance attributes or found them to be less useful. Among these, P₃ was more focused on selecting movies that satisfied the family's constraints and hence was more inclined to check the actual attributes rather than the *frequency* and *recency*. P_{6,12} speculated that the provenance attributes would be more useful if the study lasted longer or was spread across multiple sessions.

On the other hand, six participants (P_{4,8,9,13,14,15}) found utility in tracking and reviewing their provenance in real-time. P₁₃ actively used the provenance attributes to keep track of their own analysis and not become stuck while making a choice.

P₄ realized that after satisfying all family members' movie constraints, they were free to give their own recommendations. They saw they had not interacted with "Production Budget" and "Worldwide Gross" too much, so they created a scatterplot with these attributes, mapped *frequency* and *recency* on *color* and *size*, and utilized the visual scents to analyze unexplored movies and accordingly determine their final movie selections.

D. Reasons for Using the Provenance Attributes

Answer questions when unable to recall analysis behavior. Due to the study design, answering questions was the most common reason to use the provenance attributes. Reviewing analytic behavior as a quantized set of variables (*P₃*) and using them in visual encodings (*P_{1,14}*) and data transformations (*P₁*) made it easy to confidently answer questions (*P₃*).

Verify and build confidence while answering questions. Sometimes participants recalled answers from memory, but using the provenance attributes helped them verify and double-check their choices (*P_{14,16}*) and gain confidence (*P_{12,14}*).

Save time during analysis. Some participants found the *recency* and *frequency*-based data transformations and aggregations particularly useful to generate quick responses to the study questions (*P₁₁*), directly saving them a lot of time (*P₁₀*).

Increase awareness, uncover new insights, and improve exploration coverage. When selecting the remaining movies, the provenance attributes helped participants optimize their analysis by increasing awareness (*P₄*), helping them avoid revisiting the same points (*P_{9,15}*), and facilitating quick decisions (*P₁₅*). Some participants uncovered new insights (*P_{2,3}*) or unexplored data (*P₃*), and identified behavioral trends (*P₃*).

Fun and Feel Good. For some participants, the provenance attributes were easy and fun to use (*P₁₃*). The real-time changes to the colors or sizes of points made them feel good and feel like they were making progress in the task (*P₁₃*).

E. Reasons for More or Less Surprise

While answering questions, participants were surprised because (1) the system's answer did not match their recollection of their own analysis, (2) the system's answer did not match their original analytic intention, (3) the system inadequately captured the participant's focus (e.g., interactions with an aggregate visualization gave equal focus to constituent data-points, which was not agreeable to the user), or (4) the system logged interactions that were perhaps accidental (our threshold to discard mouseovers under ~ 250 ms may be low for use as a proxy to compute focus). For example, *P₂* noted, "I found [an interaction] to be a Harry Potter movie, which makes sense, but I had forgotten about having interacting with it first."

On the flip side, participants were either less surprised or not surprised at all because (1) they actively used the *recency* and *frequency* attributes during analysis and thus knew what they were doing and (2) they were able to simply remember their behaviors. *P₅* noted, "I already knew what I had focused on, but far ahead in the future, these provenance attributes might be helpful to provide reasoning for my choices when I might not explicitly remember my thought process."

F. Participant Strategies to Answer Questions

To answer questions, participants often utilized provenance attributes as a combination of encodings, transformations (sort and filter), and subsequent interactions (e.g., tooltip on hover). We observed 71 different strategies. For example, Figure 4 shows five participants' different strategies to answer T6.Q1: "How similar were your interaction patterns for 'Comedy' and 'Thriller' movies?" While *P₁* created an aggregate bar chart showing the two "Genre" categories on *x*, *frequency* on *y*, and average *recency* as *color*, *P₁₄* used a scatterplot encoding *recency* (*color*) and *frequency* (*size*), faceted by "Genre".

Co-occurrence analysis. Figure 5 shows the co-occurrence of provenance attributes (*only frequency*, *only recency*, or either) as **A** combinations of visual encodings compared to **B** data transformations. In terms of visual encodings, *x* and *y* independently were used most often followed by *fill* (*color*). Interestingly, *frequency* and/or *recency* were simultaneously mapped to both *x* and *y* a solid 36 times. Between visual encodings and data transformations, standalone visual encodings were used most often (131 times) followed by their combination with record filter (*encoding_rFilter*, 48 times) and attribute sort (*encoding_aSort*, 35 times).

G. Preferences for Utilizing the Provenance Attributes

Figure 6 shows how users mapped data and provenance attributes to different visual encodings **A**, and their general preferences for encodings compared to transformations **B**.

Visual Encodings. For provenance attributes, *x* was the most preferred encoding followed by *y*, *fill*, *size*, *text*, *fillOpacity*, *row*, *annotation*, and *column*. The encodings *shape*, *stroke*, *strokeOpacity*, and *strokeWidth* were either rarely used or not used at all. Similarly, for data attributes, *x* or *y* were the most preferred encodings. The nominal "Genre" and "Content Rating" attributes were also occasionally mapped to *fill* and *row* encodings. "Title" was predominantly mapped to the *tooltip* visual encoding, followed by *x*, *annotation*, *y*, and *text*.

Visual Encodings or Data Transformations. For provenance as well as data attributes, visual encodings were most preferred followed by filtering records and sorting attributes.

H. Recency or Frequency? What was more Useful?

Three participants (*P_{1,5,14,16}*) highlighted that frequency was more useful than *recency*. *P₁₄* mentioned that they did not rely much on *recency* for selecting movies but found *frequency* helpful in identifying movies they had visited multiple times. *P₁₁* also noted that *frequency* aided in recognizing options that were still being considered. *P₁₅* suggested that while *frequency* was more useful for the short duration of the study, *recency* might be more valuable in the long term, especially for tasks like auditing or reviewing analysis after days or weeks.

VI. DISCUSSION AND TAKEAWAYS

A. Making Provenance a Core Aspect of Analysis

Our study revealed that users can use provenance attributes to answer questions with high accuracy and confidence, while

Task6, Que1: "How similar were your interaction patterns for 'Comedy' and 'Thriller' movies? Illustrate via a visualization."

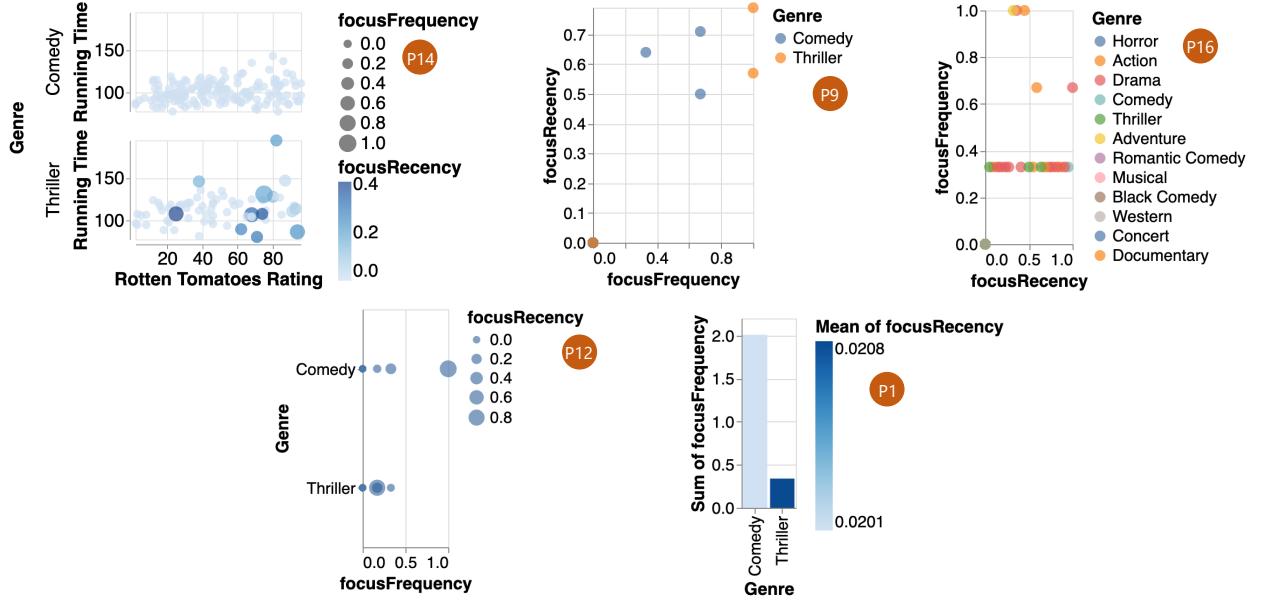


Fig. 4. Five participants' different strategies to answer the same question, T6.Q1, "How similar were your interaction patterns for 'Comedy' and 'Thriller' movies? Illustrate via a visualization." *P₁₄* created a scatterplot with "Rotten Tomatoes Rating" along x, "Running Time" along y, faceted by "Genre", colored by recency, and sized by frequency; *P₉* and *P₁₆* created a scatterplot with frequency and recency along x or y, and colored by "Genre"; *P₁₂* created a scatterplot with frequency along x, "Genre" along y, and sized by recency; *P₁* created an aggregate bar chart with "Genre" along x, total frequency along y, and colored by average recency. Except *P₁₆*, each of the above participants also applied a filter to only show movies that belong to the "Comedy" or "Thriller" genre.

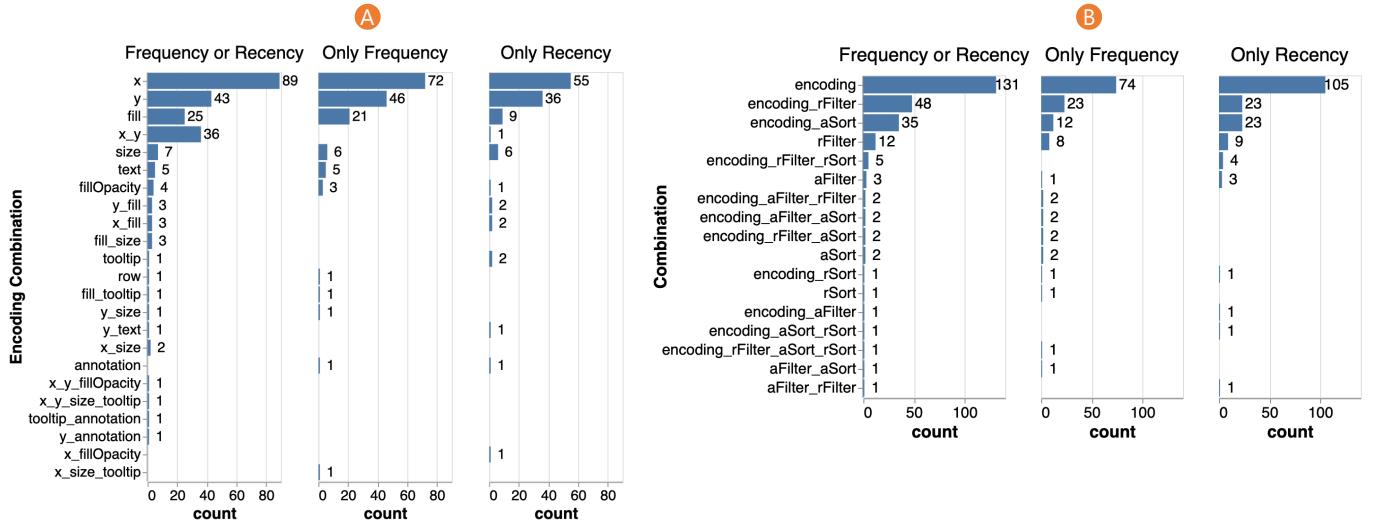


Fig. 5. Co-occurrence statistics for how users map provenance attributes (only frequency, only recency, or either) to visual encoding combinations (A), as well as general preferences for visual encodings compared to filtering, and sorting (B). Note that these statistics correspond only to the **recall** (T2, T5) and **visualize** (T3, T6) tasks; we exclude the **review** (T1) and **analyze** (T4) tasks as they were more open-ended in nature.

also sometimes being surprised. In addition, there were several instances of users requesting new capabilities not currently supported by ProvenanceLens. For example, *P₆* suggested an interesting feature to toggle between their own provenance and their mom's during analysis. *P₈* requested the ability to "undo" an accidental interaction, hinting towards an ability to directly manipulate/correct their provenance. These feature requests solidify that provenance still has unrealized utility, and should thus be considered as a core element during analysis, thereby calling for visual data analysis tools to inherently support it.

B. Fostering Provenance-driven (Not Data-driven) Analysis.

Many participants visualized provenance by mapping it to visual encodings (e.g., color) to keep track of their ongoing analytic progress (e.g., to avoid revisiting the same points). Some participants also actively sorted and filtered the data attributes and records by provenance attributes to either reduce their search space or look-up what has been previously considered. These are examples of provenance-driven (not data-driven) analysis wherein provenance information is used to steer the analysis process. Similar to how data-driven analysis focuses

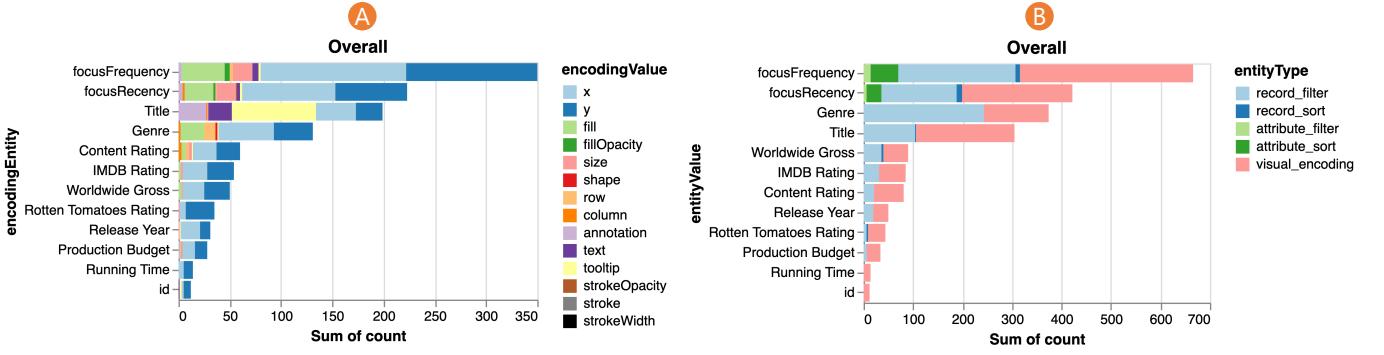


Fig. 6. User preference when mapping attributes to (A) different visual encodings and (B) using visual encodings in general compared to data transformations.

on using data to guide decision-making and insight generation, provenance-driven analysis can surface the lineage and context of data and interactions to determine next steps and ensure the accuracy, reliability, and interpretability of the analysis. We thus call for tools to support both these analysis paradigms.

C. Integrating Provenance-Tracking & Visual Data Analysis

During our study, participants often used provenance attributes together with data attributes. For example, participants mapped provenance attributes to *x*, *y*, *fill*, and *size*, which are encodings commonly used for data attributes. Furthermore, there were cases where a participant replaced a data attribute (that was mapped to a visual encoding), with a provenance attribute (and vice-versa). In other cases, participants had a preferred way of interacting with analytic provenance (e.g., always mapping provenance to *fill*) and they consistently used the same strategies during analysis or when answering questions. These behaviors suggest that tools should offer both data and provenance attributes for more flexible workflows.

D. Affording Flexibility in Mapping Provenance to Encodings

Our study revealed that users often visualized provenance on visual encodings such as *x*, *y*, and *tooltip*, something not often seen in existing tools. Furthermore, while performing tasks T2 (Recall) and T4 (Analyze), some participants used multiple visual encodings for the same question in succession, to verify the results. For example, *P₄* first mapped *frequency* to *fill*, but found it hard to accurately differentiate between different shades and hues (which can be hard for colorblind users). To verify their takeaways, they mapped *frequency* to *size* instead, but were worried about occlusion, depending on the *x/y* point positions. Consequently, they applied a double-encoding by also mapping *frequency* to *x*. This flexibility shows that if an encoding is unavailable and another is inferior, then a third encoding can still be effective, underscoring our core goal to afford flexibility in mapping provenance during analysis.

E. Comparing Provenance Encodings and Transformations

During our study, participants regularly mapped provenance to visual encodings or applied data transformations. However, depending on the question type, one technique can be more efficient than the other. Consider T6.Q2 wherein participants were tasked to show **three movies** that they most recently

interacted with. Those who applied a *recency* filter often struggled with this task due to ties in the number of times a movie was interacted with. For example, *P₁* asked, “*How do I get exactly three?*” Filtering for top-N movies (records) or movie characteristics (attributes) is less trivial as it is impossible to guess what range will produce the exact number of items. Mapping provenance to positional encodings (e.g., *x*) or filtering based on ordinal *provenance rank* instead of a quantitative *provenance score* can make it easier to spot ties.

F. Supporting Collaboration during Visual Data Analysis

In tasks T1–T3 of our user study, participants had to utilize provenance attributes to review and answer questions about another user’s (mom’s) analysis. In doing so, we indirectly studied how provenance attributes can facilitate asynchronous collaboration. By analyzing another user’s provenance in terms of what they looked at, when, and for how long, the user can not only verify or find flaws in prior analyses but also formulate a starting point or become unstuck [66]. However, overreliance on another user can hamper creativity and result in “herd behavior” [66]. Balancing these two behaviors can result in efficient collaborative analysis and decision-making.

G. Enhancing Modeling of Provenance Attributes

Our approach of modeling provenance attributes has some implications and limitations. First, mapping *recency* and *frequency* to visual encodings, such as darker and larger points for more recent or frequent data, can foster confirmation bias [72] by leading users to unknowingly and disproportionately prioritize recently or frequently interacted data points. Such behavior is an unfortunate consequence of the recency effect [73] and the frequency effect [74]. As a result, users may potentially overlook less accessed but relevant information. To mitigate this bias, a reversed provenance scale that emphasizes older and less frequent data (as darker and larger) can provide a more balanced view in decision-making processes.

Next, our approach to normalize recency and frequency values to a uniform ranked scale from zero to one, instead of displaying absolute values like raw timestamps and interaction counts, can simplify comparisons but sacrifice explainability. For instance, without timestamps, users may not understand the exact sequence or timing of interactions, i.e., whether frequent interactions occurred closely together or were spread

out over time. Additionally, users may overlook the original context (i.e., the data attributes and records themselves) or the rationale of the previous user whose provenance they are reviewing. Such lack of clarity can lead to incomplete or biased interpretations in decision-making.

Lastly, provenance attributes can be modeled in multiple ways. For *frequency*, the default *relative* strategy divides the “total interaction units” for an attribute by the maximum value among all attributes; the *absolute* strategy divides each value by the sum of all values; and the *binary* strategy treats zero interactions=0 and at least one interaction=1. For *recency*, the default *relative* strategy determines the value based on the sequence (or rank) of interactions; the *absolute* strategy determines the value based on the actual time duration between interactions; and the *binary* strategy assigns a value=1 for the most recent interaction and a value=0 for all other interactions. Our conceptualization of provenance attributes seamlessly supports all of these models for subsequent visualization.

H. Facilitating Co-Adaptive Guidance during Analysis

In this study, we studied how users can flexibly map provenance attributes to different encodings and data transformations. We also believe our system can be extended to enable co-adaptive guidance [75], wherein the system (not user) can recommend provenance mappings as orienting guidance [18] to draw the user’s attention to specific analytic behaviors such as exploration biases [3]. In response, the user can accept, reject, or override these recommendations, so that the system can learn the user’s preferences to inform future analytic discourse. In this way, users and systems can work together to provide tailored help, thereby improving the analysis processes.

I. Reflecting on Provenance-focused User-Study Designs

In our two pilot studies, participants generally ignored the provenance attributes during analysis, only using them later for subsequent question-answering. This behavior was likely due to the task’s nature or participants’ ability to simply recall their analysis. In response, we added the “review” (your mom’s) task and the “answer questions” (about your mom’s analysis) task to force participants to interact with the provenance attributes that are not modeled on their own analysis, which helped a bit. Still, during the actual study, *P₁₂* noted, “*My memory still works pretty well in the short time frame,*” and *P₆* speculated, “*Maybe if it were a longer session or spread across multiple sessions, then [provenance attributes] would be more useful.*” One key insight in designing such short-duration user studies on provenance is the importance of pilot studies as they can help refine study protocols and overcome confounds. We hope our experiences designing this user study can help other researchers better design their own studies in the future.

VII. LIMITATIONS AND FUTURE WORK

There were five main limitations of our study. First, our study (designed to be a design probe) lacked a baseline condition where participants complete tasks without interacting with provenance attributes, which would have allowed for a more rigorous comparison of performance between conditions.

Second, our study was conducted in a single session lasting under 90-minutes, which may not fully capture the long-term benefits of visualizing provenance during analysis (i.e., beyond recall-based question-answering tasks); future studies should consider longer durations to validate sustained impact.

Third, while our prototype lets users customize the range of provenance attributes (e.g., darker points can be mapped to smaller or larger values), we disabled this functionality during the study to minimize users’ cognitive load. Systematically studying this feature, i.e., if one range leads to more unique data discoveries [4] while another leads to more data revisits [3] is future work.

Fourth, we currently model focus by equally weighting all interactions but future work can weight differently, e.g., recent interactions more than the older ones [9].

Lastly, we modeled *recency* and *frequency* only based on mouse interactions such as clicks and hovers, which may not be a complete proxy for focus. In fact, *P₆* expressed some confusion, “*The wording of what ‘interaction’ means was not clear, [I] needed help to understand that.*” A similar confusion had also come up during our pilot studies, but we tried to address it by clearly explaining to users how provenance is computed, and ensuring that users become acquainted with it during the practice. We also posit some of this (recurrent) confusion may be due to participants’ surprise upon seeing the system’s projection of their provenance not matching their own expectation. However, we acknowledge that our modeling of provenance from hover interactions with data points, might be *strict* as it does not account for other forms of engagement, like focusing on specific clusters of data points (that remain after applying a filter). Still, we chose it for its practicality and because it likely reflects deliberate user engagement. Future work may utilize other mechanisms such as user gaze to more accurately model focus (e.g., which attribute(s) and data point(s) is the user actually looking more at).

In general, based on the encouraging findings of this design probe, we next plan to further understand the interplay between data and provenance attributes, particularly how users organically utilize these attributes in high-stakes analysis tasks (unlike a harmless movie selection task as in this study).

VIII. CONCLUSION

In visual data analysis systems, provenance is often logged but only reviewed after analysis, or is hardwired, limiting user access and control. In this work, we utilized provenance as an attribute *during* analysis, tracking both *recency* and *frequency* of user interactions with data. To evaluate the utility of provenance attributes, we integrated them into a prototype visual data analysis system, ProvenanceLens, which allows users to track and visualize recency and frequency by mapping them to visual encodings (e.g., color or size) and data transformations (sort or filter). An exploratory study with sixteen users found that provenance attributes can help users accurately and confidently review and answer questions about their analysis, often surprising them and facilitating self-reflection.

REFERENCES

- [1] C. North, R. Chang, A. Endert, W. Dou, R. May, B. Pike, and G. Fink, “Analytic Provenance: Process + Interaction + Insight,” in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2011, pp. 33–36.
- [2] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen, “Characterizing Provenance in Visualization and Data Analysis: An Organizational Framework of Provenance Types and Purposes,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 31–40, 2016.
- [3] A. Narechania, A. Coscia, E. Wall, and A. Endert, “Lumos: Increasing Awareness of Analytic Behavior during Visual Data Analysis,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 1009–1018, 2022.
- [4] M. Feng, C. Deng, E. M. Peck, and L. Harrison, “HindSight: Encouraging Exploration through Direct Encoding of Personal Interaction History,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 351–360, 2017.
- [5] 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, 2007.
- [6] G. A. Miller, “The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information,” *Psychological Review*, vol. 101, no. 2, p. 343, 1994.
- [7] 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, 2014.
- [8] H. Stitz, S. Gratzl, H. Piringer, T. Zichner, and M. Streit, “Knowledge Pearls: Provenance-Based Visualization Retrieval,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 120–130, 2019.
- [9] Z. Zhou, X. Wen, Y. Wang, and D. Gotz, *Modeling and Leveraging Analytic Focus During Exploratory Visual Analysis*. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: <https://doi.org/10.1145/3411764.3445674>
- [10] A. Narechania, K. Odak, M. El-Assady, and A. Endert, “ProvenanceWidgets: A Library of UI Control Elements to Track and Dynamically Overlay Analytic Provenance,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 1235–1245, 2025.
- [11] 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, 2008.
- [12] K. Xu, A. Ottley, C. Walchshofer, M. Streit, R. Chang, and J. Wenskovitch, “Survey on the Analysis of User Interactions and Visualization Provenance,” *Computer Graphics Forum*, vol. 39, no. 3, pp. 757–783, 2020.
- [13] J. Perry, C. D. Janneck, C. Umoja, and W. M. Pottenger, “Supporting Cognitive Models of Sensemaking in Analytics Systems,” *DIMACS*, 2009.
- [14] P. H. Nguyen, K. Xu, A. Bardill, B. Salman, K. Herd, and B. W. Wong, “SenseMap: Supporting Browser-based Online Sensemaking through Analytic Provenance,” in *IEEE VAST*, 2016.
- [15] K. Madanagopal, E. D. Ragan, and P. Benjamin, “Analytic Provenance in Practice: The Role of Provenance in Real-World Visualization and Data Analysis Environments,” *IEEE Computer Graphics and Applications*, vol. 39, no. 6, pp. 30–45, 2019.
- [16] Z. Bylinskii, N. W. Kim, P. O’Donovan, S. Alsheikh, S. Madan, H. Pfister, F. Durand, B. Russell, and A. Hertzmann, “Learning Visual Importance for Graphic Designs and Data Visualizations,” in *ACM UIST*, 2017.
- [17] S. Gomez and D. Laidlaw, “Modeling Task Performance for a Crowd of Users from Interaction Histories,” in *ACM CHI*, 2012.
- [18] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski, “Characterizing Guidance in Visual Analytics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 111–120, 2017.
- [19] A. Walch, M. Schwärzler, C. Luksch, E. Eisemann, and T. Gschwandtner, “LightGuider: Guiding Interactive Lighting Design using Suggestions, Provenance, and Quality Visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 569–578, 2020.
- [20] A. Endert, P. Fiaux, and C. North, “Semantic Interaction for Visual Text Analytics,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2012, pp. 473–482.
- [21] L. Bavoil, S. P. Callahan, P. J. Crossno, J. Freire, C. E. Scheidegger, C. T. Silva, and H. T. Vo, “Vistrails: Enabling Interactive Multiple-View Visualizations,” in *VIS 05. IEEE Visualization, 2005*. IEEE, 2005, pp. 135–142.
- [22] Y. B. Srinivasan, D. Gotz, and J. Lu, “Connecting the Dots in Visual Analysis,” in *IEEE VAST*, 2009.
- [23] Y. Chen, S. Barlowe, and J. Yang, “Click2Annotate: Automated Insight Externalization with Rich Semantics,” in *2010 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2010, pp. 155–162.
- [24] S. Gratzl, A. Lex, N. Gehlenborg, N. Cosgrove, and M. Streit, “From Visual Exploration to Storytelling and Back Again,” in *Computer Graphics Forum*, 2016.
- [25] E. Wall, A. Narechania, A. Coscia, J. Paden, and A. Endert, “Left, Right, and Gender: Exploring Interaction Traces to Mitigate Human Biases,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 966–975, 2022.
- [26] J. E. Block, S. Esmaeli, E. D. Ragan, J. R. Goodall, and G. D. Richardson, “The Influence of Visual Provenance Representations on Strategies in a Collaborative Hand-off Data Analysis Scenario,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 1113–1123, 2023.
- [27] A. Wattenberger, “Footsteps for VS Code,” <https://marketplace.visualstudio.com/items?itemName=Wattenberger.footsteps>, 2021, Accessed: June 15, 2024.
- [28] K. Gadhav, Z. Cutler, and A. Lex, “Persist: Persistent and Reusable Interactions in Computational Notebooks,” in *Computer Graphics Forum*, 2024.
- [29] K. Eckelt, K. Gadhav, A. Lex, and M. Streit, “Loops: Leveraging Provenance and Visualization to Support Exploratory Data Analysis in Notebooks,” *OSF Preprint*, 2023. [Online]. Available: <https://doi.org/10.31219/osf.io/79eyn>
- [30] Y. Ding, J. Wilburn, H. Shrestha, A. Ndlovu, K. Gadhav, C. Nobre, A. Lex, and L. Harrison, “reVISit: Supporting Scalable Evaluation of Interactive Visualizations,” in *IEEE VIS*, 2023.
- [31] T. Ellkvist, D. Koop, E. W. Anderson, J. Freire, and C. Silva, “Using Provenance to Support Real-Time Collaborative Design of Workflows,” in *Provenance and Annotation of Data and Processes*, 2008.
- [32] A. Savyaghad and M. Tory, “Exploiting Analysis History to Support Collaborative Data Analysis,” in *Proceedings of the 41st Graphics Interface Conference*, 2015, pp. 123–130.
- [33] S. K. Badam, Z. Zeng, E. Wall, A. Endert, and N. Elmquist, “Supporting Team-First Visual Analytics through Group Activity Representations,” in *Graphics Interface*, 2017.
- [34] “Google Analytics,” <https://marketingplatform.google.com/about/analytics>, Accessed: June 15, 2024.
- [35] “Hotjar,” <https://www.hotjar.com>, Accessed: June 15, 2024.
- [36] A. Drachen, “Behavioral Telemetry in Games User Research,” *Game User Experience Evaluation*, pp. 135–165, 2015.
- [37] T. C. Kohwalter, L. G. P. Murta, and E. W. G. Clua, “Capturing Game Telemetry with Provenance,” in *2017 16th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*. IEEE, 2017, pp. 66–75.
- [38] S. P. Callahan, J. Freire, E. Santos, C. E. Scheidegger, C. T. Silva, and H. T. Vo, “VisTrails: Visualization meets Data Management,” in *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, 2006, pp. 745–747.
- [39] W. Aigner, S. Hoffmann, and A. Rind, “EvalBench: A Software Library for Visualization Evaluation,” in *Computer Graphics Forum*, vol. 32, no. 3pt1. Wiley Online Library, 2013, pp. 41–50.
- [40] M. Okoe and R. Jianu, “GraphUnit: Evaluating Interactive Graph Visualizations Using Crowdsourcing,” in *Computer Graphics Forum*, vol. 34, no. 3. Wiley Online Library, 2015, pp. 451–460.
- [41] Z. Cutler, K. Gadhav, and A. Lex, “Trtrack: A Library for Provenance-Tracking in Web-Based Visualizations,” in *2020 IEEE Visualization Conference (VIS)*, 2020, pp. 116–120.
- [42] E. Arroyo, T. Selker, and W. Wei, “Usability tool for analysis of web designs using mouse tracks,” in *CHI’06 extended abstracts on Human factors in computing systems*, 2006, pp. 484–489.
- [43] J. Nielsen and K. Pernice, *Eyetracking web usability*. New Riders, 2010.
- [44] P. Baudisch, D. Tan, M. Collomb, D. Robbins, K. Hinckley, M. Agrawala, S. Zhao, and G. Ramos, “Phosphor: Explaining Transitions in the User Interface Using Afterglow Effects,” in *Proceedings of the 19th annual ACM symposium on User Interface Software and Technology*, 2006, pp. 169–178.

[45] J. Hill and C. Gutwin, "Awareness Support in a Groupware Widget Toolkit," in *Proceedings of the 2003 ACM International Conference on Supporting Group Work*, 2003, pp. 258–267.

[46] M. Angelini, G. Blasilli, S. Lenti, A. Palleschi, and G. Santucci, "Cross-Widgets: Enhancing Complex Data Selections through Modular Multi Attribute Selectors," in *Proceedings of the International Conference on Advanced Visual Interfaces*, 2020, pp. 1–9.

[47] C. Gutwin, "Traces: Visualizing the immediate past to support group interaction," in *Graphics interface*. Citeseer, 2002, pp. 43–50.

[48] J. Boy, F. Detienne, and J.-D. Fekete, "Storytelling in information visualizations: Does it engage users to explore data?" in *ACM CHI*, 2015.

[49] M. Feng, E. Peck, and L. Harrison, "Patterns and Pace: Quantifying Diverse Exploration Behavior with Visualizations on the Web," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 501–511, 2019.

[50] A. Ottley, R. Garnett, and R. Wan, "Follow The Clicks: Learning and Anticipating Mouse Interactions During Exploratory Data Analysis," in *Computer Graphics Forum*, vol. 38, no. 3. Wiley Online Library, 2019, pp. 41–52.

[51] D. Gotz, S. Sun, and N. Cao, "Adaptive Contextualization: Combating Bias During High-Dimensional Visualization and Data Selection," in *Proceedings of the 21st International Conference on Intelligent User Interfaces*, 2016, pp. 85–95.

[52] E. Wall, L. M. Blaha, L. Franklin, and A. Endert, "Warning, Bias May Occur: A Proposed Approach to Detecting Cognitive Bias in Interactive Visual Analytics," in *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*, 2017, pp. 104–115.

[53] A. Skopik and C. Gutwin, "Improving Revisitation in Fisheye Views with Visit Wear," in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 2005, pp. 771–780.

[54] S. Kaasten, S. Greenberg, and C. Edwards, "How People Recognise Previously Seen Web Pages from Titles, URLs and Thumbnails," in *HCI*, 2002.

[55] W. C. Hill, J. D. Hollan, D. Wroblewski, and T. McCandless, "Edit Wear and Read Wear," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1992, pp. 3–9.

[56] J. Alexander, A. Cockburn, S. Fitchett, C. Gutwin, and S. Greenberg, "Revisiting Read Wear: Analysis, Design, and Evaluation of a Footprints Scrollbar," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2009, pp. 1665–1674.

[57] W. Oliveira, L. M. Ambrósio, R. Braga, V. Ströele, J. M. David, and F. Campos, "A Framework for Provenance Analysis and Visualization," *Procedia Computer Science*, 2017.

[58] T. Kohwalter, T. Oliveira, J. Freire, E. Clua, and L. Murta, "Prov Viewer: A graph-based visualization tool for interactive exploration of provenance data," in *IPAW*, 2016.

[59] P. Chen, B. Plale, Y.-W. Cheah, D. Ghoshal, S. Jensen, and Y. Luo, "Visualization of Network Data Provenance," in *HiPC*, 2012.

[60] B. Schneiderman, "Direct manipulation: A step beyond programming languages," *Computer*, 1983.

[61] M. Wolter, B. Hentschel, I. Tedjo-Palczynski, and T. Kuhlen, "A direct manipulation interface for time navigation in scientific visualizations," in *IEEE 3DUI*, 2009.

[62] B. Kondo and C. Collins, "Dimpvis: Exploring time-varying information visualizations by direct manipulation," *IEEE Transactions on Visualization and Computer Graphics*, 2014.

[63] A. Lex, N. Gehlenborg, H. Strobelt, R. Vuillemot, and H. Pfister, "UpSet: visualization of intersecting sets," *IEEE Transactions on Visualization and Computer Graphics*, 2014.

[64] Tableau, "Tableau," 2022. [Online]. Available: <https://www.tableau.com/learn/get-started/prep>

[65] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer, "Wrangler: Interactive Visual Specification of Data Transformation Scripts," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 3363–3372.

[66] A. Narechania, F. Du, A. R. Sinha, R. Rossi, J. Hoffswell, S. Guo, E. Koh, S. B. Navathe, and A. Endert, "DataPilot: Utilizing Quality and Usage Information for Subset Selection during Visual Data Preparation," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023.

[67] A. Narechania, S. Chakraborty, S. Agarwal, A. R. Sinha, R. A. Rossi, F. Du, J. Hoffswell, S. Guo, E. Koh, A. Endert, and S. Navathe, "DataCockpit: A Toolkit for Data Lake Navigation and Monitoring Utilizing Quality and Usage Information," in *2023 IEEE International Conference on Big Data (BigData)*, 2023, pp. 5305–5310.

[68] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-Lite: A Grammar of Interactive Graphics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 341–350, 2017.

[69] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman, "The Science of Visual Data Communication: What Works," *Psychological Science in the Public Interest*, vol. 22, no. 3, pp. 110–161, 2021. [Online]. Available: <https://doi.org/10.1177/15291006211051956>

[70] C. Nobre, D. Wootton, Z. Cutler, L. Harrison, H. Pfister, and A. Lex, "reVISit: Looking under the hood of interactive visualization studies," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–13.

[71] J. Brooke, "SUS: A Quick and Dirty Usability Scale," *Usability Evaluation in Industry*, vol. 189, no. 194, pp. 4–7, 1996.

[72] R. S. Nickerson, "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, 1998.

[73] R. C. Atkinson and R. M. Shiffrin, "Human Memory: A Proposed System and its Control Processes," in *Psychology of Learning and Motivation*, 1968.

[74] A. Tversky and D. Kahneman, "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology*, 1973.

[75] F. Sperrle, A. Jeitler, J. Bernard, D. Keim, and M. El-Assady, "Co-Adaptive Visual Data Analysis and Guidance Processes," *Computers & Graphics*, vol. 100, pp. 93–105, 2021.



Arpit Narechania is an assistant professor in the Department of Computer Science and Engineering at The Hong Kong University of Science and Technology (HKUST). He directs the DataVisards group, which explores novel data experiences for human-AI collaboration for visual analytics.



Shunyan Guo is a research scientist at Adobe. Her research interest lies in the area of information visualization and visual analytics, with a focus on designing and developing visual analytics prototypes that communicate insights derived from complex data models.



Eunyee Koh is a principal scientist at Adobe. Her research interest lies at the intersection of human-computer interaction, AI/machine learning, and data science. She is passionate about how to personalize information for people.



Alex Endert is an associate professor in the School of Interactive Computing, Georgia Tech. He directs the Visual Analytics Lab, which explores novel user interaction techniques for visual analytics.



Jane Hoffswell is a research scientist at Adobe, specializing in visualization and human-computer interaction. Her research explores the design of responsive visualizations, interactive systems, and data-driven visual storytelling.