Serendipitous Explanations: Interaction-Triggered Comprehension Aids in Visualization

M. Rezaie[1], M. Tory[2], S. Carpendale[1]

[1]Simon Fraser University, Burnaby, BC, Canada
[2]Northeastern University, Boston, MA, USA

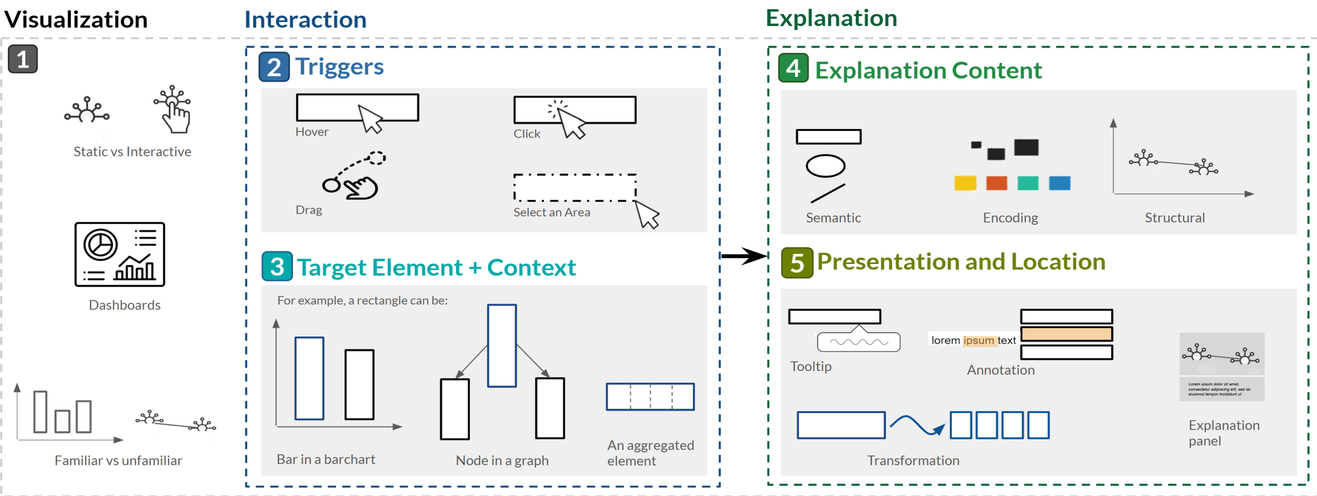

Figure 1: The Serendipitous Explanations Approach and its core dimensions with some examples.

Abstract

Evidence continues to accrue around the difficulties people have understanding new and complex visualizations, which in turn provides continued incentive to explore additional methods of supporting visualization viewers. Through leveraging viewers' spontaneous visualization sensemaking activities, we introduce a Serendipitous Explanation Approach (SEA). A significant part of SEA's contribution is making active use of a viewer's spontaneous, exploration interactions to offer in situ explanations. SEA adds semantic and structural explanations, making use of visual transformations, and additional representations of visualization elements themselves to communicate meaning through highlights, reconfigurations, and in-situ annotations. We designed and studied VisTips as an instantiation of SEA. Our study demonstrates the prevalence of spontaneous interaction and appreciation of this approach. Among the explanation types, visual transformations stood out as especially impactful. Our findings also offer practical insights for future use of the ideas in SEA, opening possibilities for seamlessly integrating explanatory visuals that support viewers' natural visualization sensemaking.

CCS Concepts

• ***Human-centered computing*** → ***Visualization theory, concepts and paradigms;*** *Empirical studies in visualization;*

1. Introduction

Interactive visualizations are powerful tools for exploring and understanding complex data, but they are often challenging for visualization viewers to comprehend, especially when encountering unfamiliar or complex visual representations [LKH*15, MMF19]. Prior studies have examined the reasons behind this difficulty [RTC24, WCS*22], and proposed various strategies to improve visualization comprehension. One common approach is the use of structured guidance, such as video-based tutorials [DWH*22], interactive step-by-step guides and scrollytelling [HS25, SWP*22] which guide viewers through a visualization's features in a structured, often linear fashion. While effective in many ways, these approaches may not accommodate viewers' natural unstructured behavior, as they often assume a linear learning process [SB19, DHF*24, RTC24].

In practice, when viewers first encounter an unfamiliar visualization, they tend to engage in interactions with visualization elements in an unstructured, trial-based manner, attempting to uncover the meaning of the representation through spontaneous exploration [RTC24, LKH*15]. This behavior, which we refer to as *spontaneous interaction*, aligns with serendipitous discovery [THC12] and principles from constructivist learning theory, which states that knowledge is built through active engagement rather than passive instruction [Hei91]. Despite its intuitive nature and prevalence, spontaneous interaction has received limited attention in visualization comprehension, leaving a substantial opportunity to leverage it for enhanced viewer onboarding.

In this paper, we introduce the ***Serendipitous Explanations Approach (SEA)***, an interaction paradigm that treats spontaneous viewer interactions as valuable signals about what the viewer wants to understand in that moment and as opportunities to deliver lightweight, in-situ explanations, embedded in the visualization itself. Rather than replacing structured onboarding methods, this paradigm complements and can operate in parallel with them by responding to viewers' natural curiosity, offering contextual explanations at the time and place of interaction. This approach empowers viewers to construct their own learning path, shaped not by scripted tutorials, but by where their eyes land and their curiosity leads. It transforms onboarding from something dictated to something discovered, allowing explanation to emerge naturally from exploration.

What sets SEA apart is twofold. First, it is grounded in the novel idea of treating spontaneous, unstructured casual probing as an opportunity for providing explanations on demand. Unlike prior onboarding strategies that rely on predefined sequences, SEA surfaces explanation when and where comprehension demand arises. While lightweight interaction techniques like tooltips and brushing-and-linking are common, they are typically used to expose raw data or visual correspondences. In contrast, SEA treats these moments of interaction as entry points for explanation, deliberately responding to viewers' sensemaking process rather than just revealing data values. Second, it expands the idea of what an "explanation" can be in a visualization context. While static tooltips or annotations, have been explored individually, SEA is novel in how it integrates these into a viewer-driven onboarding paradigm, grounded in constructivist learning theory and interaction-as-signal logic. SEA introduces interactive visual explanations, transformations such as expansions, highlights, in-situ annotations, and visual linking, that surface explanatory meaning through the visualization itself.

To explore this paradigm, we present ***VisTips***, a working instantiation of SEA demonstrated within a complex traditional visualization. VisTips provides transient explanations in response to spontaneous hover interactions, helping users make sense of an unfamiliar visualization and its visual encodings without requiring a predefined learning sequence. VisTips reuses and extends concepts like visual transformations, brushing and linking, etc.

We conducted a between-subject study comparing VisTips with a version that included a standard legend for traditional onboarding. This study was designed to validate that spontaneous viewer interactions can be harnessed for explanation, rather than to establish VisTips' superiority over other onboarding methods such as interactive step-by-step guides. Our analysis shows VisTips expedites visualization comprehension, relative to this baseline. Our findings offer insights into how viewers engage in spontaneous interaction, and how supporting this behavior through VisTips can reshape the sensemaking process. They also highlight key considerations for leveraging the approach in the design of future instantiations.

Our main contributions are: 1) *Serendipitous Explanations Approach* (SEA) that treats spontaneous viewer interactions as opportunities to deliver explanations. 2) Design and evaluation of *VisTips*, an instantiation of SEA that augments an existing visualization with light-weight, transient interaction-triggered explanations. 3) Practical insights about SEA and VisTips offering design considerations for applying these ideas to other visualizations.

2. Related Work

We explore related work in four areas: challenges of visualization comprehension, existing strategies for guided learning, foundations of serendipitous interactions and leveraging spontaneous interaction for visualization comprehension.

2.1. Visualization Comprehension

Improving visualization comprehension requires a deeper understanding of why viewers struggle with unfamiliar visualizations. Lee et al. [LKH*15] introduced NOVIS, a framework that helps explain the process of how people make sense of unfamiliar visualizations, showing that difficulties arise when viewers fail to form a coherent mental model of what is represented and how it is encoded. Other works teased apart this process to extract what makes visualizations, including primitive [WCS*22] and unfamiliar charts [RTC24, MMF19], hard to understand such as decoding the visual encoding, matching textual explanations with visual elements, and understanding relationships between elements. Building on prior understanding of viewers' needs and their natural sensemaking behaviors, we propose SEA and provide empirical insights into how interaction-triggered, in-situ transient explanations influence the sensemaking process. In Section 4.1, we identify specific comprehension needs that VisTips addresses through our approach.

2.2. Strategies for Guided Learning

Studies have investigated how video and game-based tutorials can help teach people how to read unfamiliar visualizations [SWP*22, GWL*19, TGF13]. When studying parallel-coordinates, Kwon et al. [KL16] found that video with interaction worked best. In adapting onboarding techniques for new viewers, Stoiber et al. [SGP*19] proposed a design space for visualization onboarding and compared methods such as video-based, interactive step-by-step, and storytelling tutorials [SWG*21] for teaching parallel-coordinates. Similarly, they examined onboarding for treemaps [SGA22].

Dhanoa et al. [DHF*24] proposed an interactive tour that allows viewers to navigate freely within the tour components. This approach provides some autonomy, but still requires viewer to progress through a sequential learning process. Chundury et al.

[CYC*23] integrated in-situ explanations, along with other help seeking modes such as help topic listing and guided tour; they found that in-situ explanations (point & learn) were the most useful. These approaches often rely on structured, step-by-step help mechanisms. SEA complements and extends these methods by supporting spontaneous, unstructured behaviors grounded in viewers' natural exploration patterns. SEA does not replace tutorials or onboarding, but instead offers additional support opportunities.

2.3. Foundations of Serendipitous Explanations Approach

When viewers first encounter an unfamiliar visualization, their engagement is often unstructured and exploratory. Rather than following a predefined learning path, they hover over elements, click on data points, zoom in and out, and toggle settings, not to complete a specific task, but to see how the system responds. These spontaneous, trial-based interactions help viewers get a sense of the visualization's structure, content, and interactive possibilities [RTC24].

This behavior aligns with learning by exploration [dMvO96] and constructivist learning theory [Hei91], which state that some individuals learn best by actively engaging with their environment, forming knowledge through direct experience rather than passive instruction. This suggests that viewers construct mental models of the visualization through direct interaction with visual elements, and observation of system responses. With this strategy viewers form some understandings of the visual representations and iteratively refine their understanding [LKH*15]. These spontaneous interactions serve as meaningful moments of sensemaking that can be supported rather than ignored.

Although spontaneous interaction is common and cognitively meaningful, it remains underexplored as a mechanism for visualization sensemaking. In HCI and software learning research, similar behaviors have been described as trial-and-error, or *"trying different approaches to solve a problem (trials), discarding failures (errors), and repeating until one is successful."* [MVFM22, p. 2] Trial-and-error is a preferred strategy for learning complex systems, despite the availability of structured tutorials [MVFM22]. Unlike trial-and-error in feature-rich software applications, which often involves finding efficient ways to complete tasks [MAC*20], viewers of a visualization are not necessarily searching for a solution but are instead trying to understand what is available and how the visualization functions. Several studies report spontaneous mouse movement [RTC24, BOZ*14, BE14], while others describe intentional but unstructured interactions aimed at forming mental models, for instance, by interacting with different elements to observe how they affect the visualization (e.g., highlighting, filtering, zooming, or reorganization of visual elements) [LKH*15, YEB16, PWM10].

These behaviors are also reported in early stages of exploratory visual analysis, where viewers engage in bottom-up visual discovery before developing specific goals [RJPL16, PWM10, DRRD12]. It is important to note that in these studies, participants are typically familiarized with visualization encodings through tutorials. While these tutorials may cover some interactions, viewers may still engage in exploratory interactions to uncover additional capabilities, beyond merely seeking insights from the data. Nevertheless, spontaneous interaction with visual elements remains, highlighting the importance of better leveraging such interaction as an integral part of the visualization experience.

Spontaneous interaction closely relates to serendipitous discovery which is the process of encountering unexpected yet meaningful insights during open-ended exploration [THC12]. It can be nurtured by traits like curiosity, open-mindedness, domain knowledge, and observational skill [THC12]. Serendipity emphasizes discovering insights in the data itself, while we extend this notion to uncover meaning at just the right moment. By repurposing known serendipity-supporting mechanisms toward comprehension, such as curiosity-invoking displays [TAM17] and multiple visual access points [THC12], Serendipitous Explanations opens a new space: explanation that emerges naturally from exploration.

2.4. Leveraging Spontaneous Interactions for Explanations

Taken together, this body of work highlights spontaneous interaction as a rich, underutilized part of the visualization experience, not only for exploration, but as a potential entry point for explanation. Indications suggest that supporting initial exploratory behavior can influence the sensemaking process. For example, visualizations that provide hover-based explanations encourage more exploratory, pattern-seeking activity [MTW*12], while highly interactive interfaces offering on-demand detail foster deeper viewer engagement [Mar06]. However, while prior work recognizes the value of such interactions, it has not actively leveraged these behaviors as implicit signals for delivering explanation. To the best of our knowledge, SEA is the first paradigm to leverage viewers' natural behavior specifically for visualization onboarding purposes.

3. Serendipitous Explanations Approach: Core Constructs

A serendipitous explanation is an interaction-triggered, in-situ, lightweight and transient explanatory response that emerges from a viewer's spontaneous exploration behavior and is delivered at the moment and location of emerging comprehension need. Crucially, it is embedded within and mediated through the visual representation itself; often through transformations, highlights, or reconfigurations, rather than imposed as an external instructional layer, and does not rely on a predefined learning sequence. Beyond defining a particular explanatory response, SEA introduces an interaction paradigm for onboarding. It is generative in that it reframes spontaneous interaction as an opportunity for situated explanation, suggesting new explanatory possibilities embedded within the visual representation itself wherever exploratory interaction occurs.
To operationalize SEA as an interaction paradigm, we articulate five core constructs that structure how serendipitous explanations emerge within interactive visualizations. In this description, the colored numbers like 2 refer to the overview in Figure 1.

3.1. 1 Visualization

Visualizations encode data into structured elements that also serve as cues for viewer interaction. They vary in forms (e.g., networks, maps), interaction profiles, and whether they are static or interactive, single chart or dashboard-based, and familiar or unfamiliar to the viewer. These characteristics define which elements are likely to require explanation; for example, in network diagrams, viewers often want to understand what each node represents and what the links mean [RTC24]. In dashboards, viewers not only need to interpret each chart individually but also understand how a shared data entity is represented across multiple views [RTC24]. These characteristics also shape how and where the explanations can be presented. For instance, in small multiple visualizations, where many small charts are shown in a grid, screen space is limited and redundancy is high. In such cases, explanations may need to be aggregated into a shared legend, rather than shown as explanatory bubbles in each view. Available interaction affordances also shape how explanations can be triggered. For instance, when a viewer scrolls quickly up and down a long list without pausing on any item, an explanation about the list, such as a brief note clarifying what the list represents or how it is organized, can be triggered.

3.2. Interaction

3.2.1. 2 Triggers

Interaction triggers (e.g. hovering, clicking, or scrolling) specify how moments of explanation are initiated. These actions signal which specific part of the display holds the viewer's attention and may require explanation. Triggers initiate the explanation process without an explicit query, offering a lightweight and intuitive way for the system to determine when, where, and for which element to provide or remove an explanation.

3.2.2. 3 Target Elements

Visual elements include graphical marks such as bars, nodes, lines, or regions, which convey meaning through visual channels like position, color, or shape. Textual elements such as axis labels, tooltips, or explanatory paragraphs provide linguistic context; they can be dense, ambiguous, or domain-specific, prompting viewers to seek clarification. In SEA, the elements that viewers interact with can become the focus of explanation, with explanatory content embedded directly within them through transformation, reconfiguration, etc.

3.2.3. 3 Context

Context refers to the surrounding structural, visual, or semantic conditions that influence how an element should be interpreted. For example, the same visual form (e.g., a rectangle) may represent a node in a network or a data bar in a chart. Context includes factors such as *density* of elements (that affects visual saliency and available space), as well as the *view state*, such as zoom level or filtering, each shaping what kind of explanation is most appropriate. Together, these contextual details allow the system to adapt explanations to the viewer's visual and analytical circumstances.

3.3. Explanation

3.3.1. 4 Explanation Type

The response type defines the content of the explanation provided by the system in response to a viewer's interaction. It is determined based on the type of element selected and its context, and addresses what the viewer is likely trying to understand. Common types of explanation content include **semantic explanations** (e.g., what an element represents), **structural explanations** (e.g., how it is connected or positioned within the visualization), and **encoding-based explanations** (e.g., what a visual variable like size or color represent). These types differ not only in the kind of information they convey, but also in the *granularity* and *focus* of that information. For instance, an explanation may describe a general relationship ("This line connects an author to a book") or a specific instance ("Mary Cooper published A New Grammar in 1750"). Distinguishing explanation types allows the system to selectively offer explanations that are most relevant to the viewer's momentary focus.

Tips can be generated in a variety of ways depending on the visualization context and available data. The data serves as the source of the explanation content, whether it is statically defined or dynamically assembled at runtime. A straightforward approach is to hardcode specific explanations for known elements. A more flexible variation uses fixed templates where key phrases or values (e.g., names, roles, years) are dynamically filled in from the dataset. Alternatively, tips can be dynamically generated using visualization metadata; for example, extracting encodings from the underlying specification (e.g., D3 structures). This process can be formalized within a rule-based system, where viewer interactions and element properties are matched against predefined conditions that specify what type of explanation to show.

3.3.2. 5 Presentation and Location

Response presentation refers to the form, layout, and placement through which explanations are surfaced in the visualization which can depend on screen space, interaction type, and context. Typical forms include tooltips (transient popups near the target element), side panels (persistent areas for rich detail), inline annotations (labels overlaid on the element), and callouts which use visual cues like arrows, bounding boxes, or connectors to associate explanatory content with specific targets. Additionally, the system may employ highlights, such as changes in color, opacity, or border thickness, to draw attention to related elements or structures in the visualization.

These constructs are interdependent rather than modular components. Interaction provides the initiating moment for explanation, target elements determine where explanation is anchored, context shapes how the interaction should be interpreted, and explanation type and presentation define how meaning is embedded. Together, they structure how serendipitous explanations emerge.

4. VisTips: A System Instantiating the Paradigm

To explore and instantiate the Serendipitous Explanations Approach, we developed VisTips, an interactive system that provides in-situ, on-demand explanations in response to users' spontaneous interactions such as hovering and clicking. Rather than guiding users through a predefined tutorial, VisTips leverages users' natural exploratory behavior to surface explanations precisely when and where they are needed. In building VisTips, we grounded our implementation in a set of well-documented comprehension challenges from prior visualization research [LKH*15, RTC24, MMF19]. In the following section, we outline these comprehension challenges and describe how VisTips responds to each through interaction-triggered explanatory features.

4.1. Comprehension Needs

We identify six core comprehension needs that arise during the process of making sense of a visualization. These needs, drawn from prior empirical studies [RTC24, MMF19, LKH*15], form the foundation for determining what should be explained in response to user interactions. With VisTips, our goal was to address these needs:

- **CN1: Identifying Visual Marks:** Users need to determine what

each visual mark represents on two levels: (a) its semantic identity: the data entity it refers to (e.g., a person, title, or publisher), and (b) its graphical identity: the kind of visual role it plays (e.g., a node in a network).
- **CN2: Decoding Visual Encodings:** Users need to interpret how data attributes are mapped to visual channels such as color, size, or position. This includes both standard and novel encodings.
- **CN3: Recognizing Semantic Relationships:** Users often seek to understand relationships between entities such as connections between people, organizations, events, or outcomes. These are often implicit or visually subtle.
- **CN4: Linking Textual and Visual Elements:** Users need to map textual references (labels, tips, annotations) to visual marks and vice versa. This bi-directional link is critical for interpreting explanations or narratives.
- **CN5: Receiving Support During Freeform Exploration:** Users rarely follow a fixed path. During unguided exploration, lightweight, contextual feedback is needed to support evolving sensemaking without disrupting flow.

To examine how SEA could address these needs in practice, we developed VisTips as a proof-of-concept system. Rather than designing a visualization from scratch, we augmented an existing one.

4.2. WPHP Visualization

We explored VisTips in the context of an existing visualization, WPHPVis, initially developed with and for domain experts [TRL*24]. Their dataset, Women's Print History Project (WPHP), is publicly available [Wom]. We selected WPHPVis because it incorporates a diverse range of visual channels and mark types [Mun14] while embedding meaningful semantic relationships between visual elements. In addition, WPHPVis was vetted by domain experts and a user study found that participants appreciated the design but found it complex to interpret. Using WPHPVis as a baseline ensures that VisTips is tested with an well-designed yet cognitively demanding visualization, allowing us to assess VisTips' effectiveness in improving comprehensibility. The visualization shows a network of 552 people, 1792 titles, and 798 firms, connected through 2372 person–title links and 3918 title-firm links. People and firms had different roles, such as author, printer, or publisher. VisTips augments WPHPVis with interaction-triggered explanations aimed at improving understandability.

4.3. VisTips Features

We designed a set of features in VisTips that instantiate dimensions of SEA to address the comprehension needs identified earlier. For each feature, we specify the dimensions we explained in Section 3: 2 **Trigger**, 3 **Target Element+ Context**, 4 **Explanation Content**, 5 **Presentation and Location**. Because the first dimension is visualization itself, our numbering starts at 2.

F1: Progressive Exploration of Aggregates and Elements: When users 2 **hover** over an 3 **aggregated visual mark** (e.g., a bar representing multiple people under the *Person* column), 5 **the visualization transforms in-place** by expanding the mark into discrete, individual components (Figure 2a, yellow squares), revealing its underlying structure. Each resulting square can then be individually **hovered** to surface a 5 **tooltip** that provides explanations about the 4 **specific entity (e.g., a person) and its relationship (e.g., a link to a title)**. For example, hovering over a person square displays an explanation such as: "*[Author Name]* contributed as *[role]* to *[Title]* in *[Year]*." This layered interaction supports **CN1a** and **CN1b** by disambiguating both what the mark refers to and how it is visually encoded. It also addresses **CN3** by surfacing contextual relationships between entities.

F2: Annotated Connections Between labels and Visuals: 2 **hovering** over 3 **textual labels such as a column headings** triggers 5 **in-situ annotations and highlighting** that clarify 4 **structural relationships between visual elements** (Figure 2b). This supports CN3 by making the connections between entities explicit, and CN5 by offering lightweight help during casual probing.

F3: Interactive Legend: Miniature Explanation Environment: To support visualization viewers who begin their understanding with the legend, we implemented an interactive legend that consists of a miniature version of the visualization, structured into two distinct rows: the first row introduces **individual visual elements** (e.g., people, titles, firms), while the second row demonstrates **possible interaction behaviors** associated with those elements (e.g., highlighting, expansion). In line with SEA, explanations are triggered through interaction and are not presented as a fixed instructional sequence. When users 2 **hover** over 3 **elements in the legend**, the system displays 5 **tooltips and lightweight annotations** that clarify how to read and interpret the visual representation. These explanations cover 4 **the semantic identity of the element** (e.g., whether it represents a person or a title) as well as its 4 **graphical encoding** (e.g., color). For example, hovering over a line between a person and a title displays: "This line represents a contribution link; it shows that the person participated in the making of the title.Color encodes the person's role, such as publishing or authoring" (Figure 2c). This supports users in learning how to decode both the data meaning and the visual grammar of the system, without requiring trial-and-error within the main view. This design supports **CN1a**, **CN1b**, **CN2** and **CN3**).

F4: Text-to-Visuals Mapping: To support users in navigating the often difficult task of connecting 3 **textual references** (e.g., descriptions, annotations, or labels) to their corresponding 3 **visual elements** in a complex visualization, we implemented 5 **visual highlighting**. When viewers 2 **hover** over a word or phrase in the interface, the system highlights the associated visual marks in the chart, such as nodes, bars, or links. This establishes an immediate visual connection between narrative content and graphical representation clarifying 4 **structural relationships between text and visuals** (Figure 2d). This feature directly addresses **CN4** and **CN5**.

F5: Hoverable Buttons as a Dropdown Alternative: Instead of relying on a dropdown menu to filter by roles, we present all roles (e.g., Author, Publisher, Printer) as visible, 3 **color-coded buttons**. When viewers 2 **hover** over a role button, the system 5 **highlights the corresponding links and nodes** in the main visualization that are associated with that role (Figure 2e). This interaction provides an immediate visual response, allowing users to quickly see 4 **how entities such as people and firms are linked to titles through a given role**. Also, the system displays a 5 **textual explanation** that states: "People and firms had different roles in the process of publishing titles." This combination of visual highlighting and textual messaging helps viewers understand that color encodes

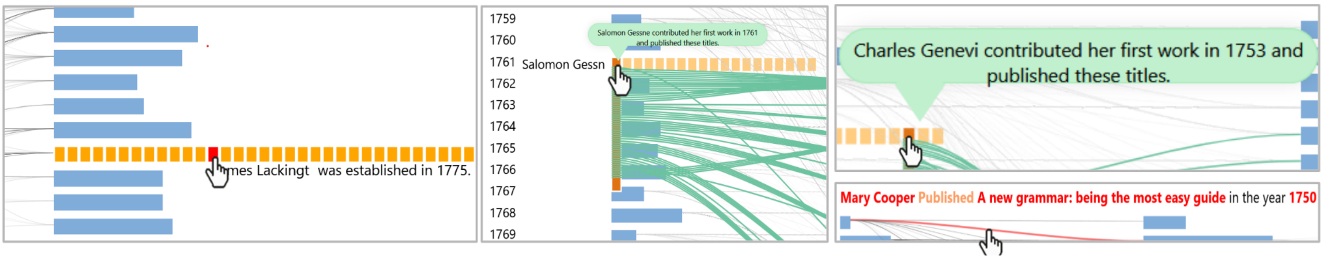

(a) Contextual explanations for individual visual elements.

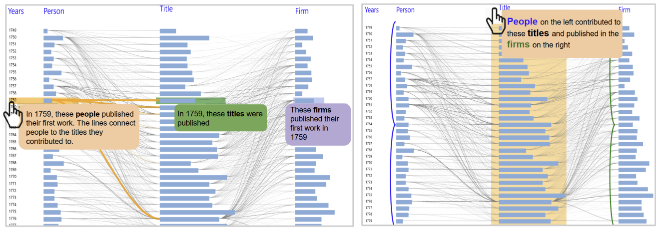

(b) Annotated connections between labels and related elements.

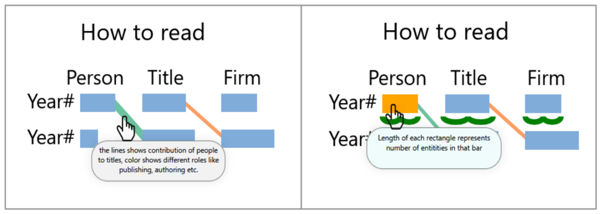

(c) Interactive legend with visual grammar and interaction cues.

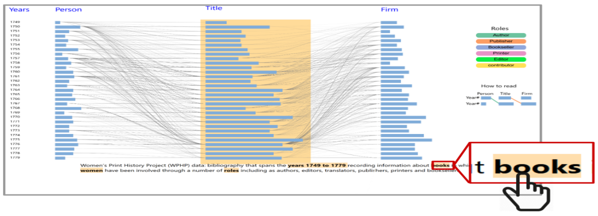

(d) Highlighting to link textual and visual elements.

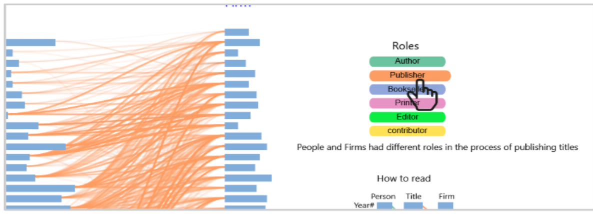

(e) Hovering over role buttons instantly highlights related links and provides a description, facilitating quick understanding of connections.

Figure 2: Six VisTips features implementing SEA. Each subfigure illustrates a different type of interaction-triggered explanation.

the 'role' attribute, and reflects a specific kind of contribution such as authoring, printing, publishing (CN2, CN3).

4.4. Implementation Details

Interaction-triggered explanations were defined through explicit event bindings on visualization elements (e.g., mouseover, mouseout, click). For each interactive element type (e.g., aggregated bar, link, role button), explanation behavior was hard-coded using conditional logic that specified: (1) the trigger event, (2) the explanation such as text template, visual transformation, highlight etc. VisTips was implemented with D3.js [BOH11].

5. Studying Serendipitous Explanations Through VisTips

We conducted a study to explore VisTips as a proof-of-concept implementation of SEA. Rather than evaluating it against existing tutorial or structured onboarding systems, (e.g. step-by-step guides) our goal was to examine whether leveraging viewers' natural interactions to trigger lightweight, in-situ explanations improves comprehension of an unfamiliar visualization. We therefore compared two versions of WPHPVis: the original [TRL*24] which includes standard affordances such as data tooltips and a static legend and an enhanced version that extends this baseline by embedding VisTips features. By comparing these two variants of the same visualization, we aimed to gain insights into how the presence of VisTip features influenced viewers' learning strategies and sensemaking processes. We used a between-participants design, with random assignment to avoid learning [LFH17] and order effects, such as practice effects and boredom [Fie13]. The study consisted of two phases. First, participants were asked to explore the visualization freely until they felt they had gained an understanding of it, without any tutorial. This phase allowed us to observe how leveraging serendipitous interaction contributed to their sensemaking process. Next, participants completed a series of tasks designed to assess their understanding and experience. We measured task completion time and collected subjective ratings.

5.1. Participants

We recruited 24 participants through fliers and word of mouth, from diverse academic and professional backgrounds. Participants ranged between 18 and 64 years old, with the majority (14/24) in the 25-34 range. 13 identified as female, 10 as male, and 1 preferred not to disclose. To capture a broad spectrum of perspectives, our participants consisted of 16 students (grad and ugrad), and 8 professionals from various disciplines, including Computing Science (6), Architecture (2), Psychology (1), Criminology (1), Business (1), Chemistry (1), and Statistics (1), among others.

Participants had varying levels of experience with data visualization. 11 participants reported encountering visualizations in media but had never formally studied them, 6 participants actively followed visualization-related topics but lacked hands-on experience. Three participants had been creating visualizations for at least a year, while 4 participants had no prior visualization knowledge. None of the participants had seen or used WPHPVis before. This diverse participant pool allowed us to explore insights into VisTips' effectiveness across varying expertise levels.

5.2. Tasks

Participants completed six tasks arranged in order of increasing complexity. We used a fixed task order for all participants—starting with simpler tasks and gradually introducing more complex ones to ensure that participants first developed an understanding of the visualization before engaging with more cognitively demanding tasks. By maintaining a structured progression of tasks from easy to complex, we aimed to mitigate varying learning effects and to have more reliable comparisons. It also allowed us to assess how VisTips supports sensemaking across a range of task demands, from lookup to complex analytical reasoning. The tasks were:

Easy: Two tasks required participants to extract information from a single visual element, needing only one interaction to complete.

Intermediate: Two tasks required participants to interpret multiple visual encodings to understand different aspects of the data.

Complex: This task required participants to synthesize information from multiple data attributes, recognizing how they were interconnected within the given context.

Trend-Identification: This task asked participants to recognize patterns or trends by analyzing multiple relationships.

5.3. Procedure

After a demographic questionnaire, the experimenter explained the process of the study. The eye-tracker was then calibrated to accurately track eyegaze. Participants began the sensemaking process without any tutorial. No separate familiarization phase was provided; participants' initial free exploration served as their only exposure to the visualization prior to the tasks. They were instructed to interact with the visualization naturally, and were informed that the visualization was interactive. They were encouraged to think aloud as they attempted to understand the visualization.

Once participants indicated that they fully understood the visualization, they proceeded to the tasks (Section 5.2). There was no time limit. Interaction logs, gaze data, task completion times, and subjective ratings were recorded. Additionally, the entire session was audio-video recorded to capture verbal responses and interaction behaviors. After completing the tasks, participants were given a Likert-scale questionnaire to assess their perceptions of the visualization, ease of use, and overall experience.

For participants in the VisTips condition, an additional semi-structured interview focused on their experiences with VisTips, specific interaction behaviors, comprehension strategies, and usability concerns was conducted. Participants also completed an additional set of Likert-scale questions for each VisTips feature, evaluating their usefulness in supporting sensemaking.

The study concluded with a debriefing and exit interview session where participants were shown the version they had not seen and invited to provide additional feedback or ask questions about the study. All sessions lasted approximately 45 minutes, and participants were appreciated with a $25 gift card for their time.

5.4. Data analysis

We analyzed the data using a combination of statistical tests and an inductive analysis [CS90]. Qualitative data was derived from video recordings that included eye-tracking data, verbalizations, interviews, and log interaction data. Following the guidelines of inductive analysis [VG22], the primary researcher performed an open coding pass on the data. Through discussions with other two members of our research team, we identified recurring themes around how participants made use of VisTips, teasing apart exploratory behavior in sensemaking, and potential areas of improvement and alternatives in leveraging this behavior. Task completion times and participants' subjective ratings on usability, engagement, and perceived comprehension were analyzed using the Mann-Whitney U test, as both task time distributions and ordinal ratings did not meet the assumptions of normality or equal intervals [BS99].

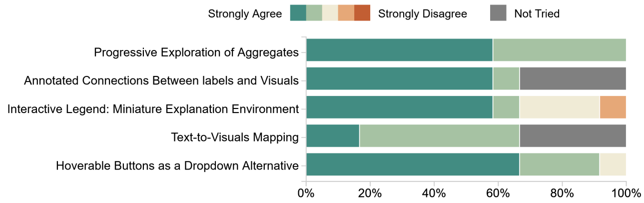

Figure 3: Participants' ratings on usefulness of VisTips features.

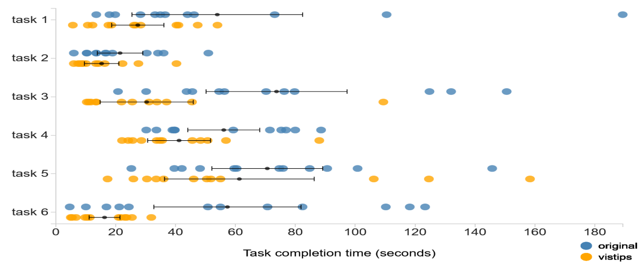

Figure 4: Breakdown of results by task and condition. Colored dots represent task completion times, with black dots indicating the mean values and bars representing 95% CIs.

6. Findings

Findings are presented as qualitative themes that characterize how spontaneous interaction unfolds and how interaction-triggered explanations shape sensemaking. Such an interpretive, exploratory stance is particularly appropriate for investigating emergent interaction phenomena aimed at advancing design concepts [MM25]. Our findings support the central premise of SEA: that viewers naturally engage in spontaneous, unstructured interaction when encountering unfamiliar visualizations, and that these moments can be leveraged to deliver lightweight, in-situ explanations. Participants reported highly positive experiences with VisTips, which instantiated the Serendipitous Explanations approach through interaction-triggered explanations. In the sections that follow, we first describe participants' overall impressions and quantitative outcomes using VisTips. We then examine which VisTips features were perceived as most effective, before turning to a deeper analysis of how SEA reshaped sensemaking and exploration, and how future systems might better support or extend this interaction-driven approach.

All 12 participants in the VisTips condition were strongly positive about their experience. During the post-task semi-structured interviews, participants highlighted how VisTips accelerated their understanding by providing on-demand explanations and easy interactions. For instance, P17 said, *"It is surprising that when you hover on anything, it explains, makes things quick to understand"*. VisTips was easy to use according to P3 who said *"It is much easier to interact with than the traditional ones ['ones' refers to visualizations they had seen before]"*.

We asked VisTips participants to evaluate each VisTips feature and share opinions on their helpfulness. They completed Likert scale questions assessing whether each specific feature aided their understanding of the visualization. As shown in Figure 3, participants found most features helpful in enhancing their understanding of the visualization. Participants reported the highest level of agreement for the Progressive Exploration (F1) feature, indicating that it was perceived as the most helpful; After interacting with the bar, P1 said, *"oh, it breaks apart, cool, I can interact with each discrete data.[...] It is like a drawer opening up to reveal what is inside"*. The Hoverable Buttons (F5) were used the most and were helpful in the sensemaking process. P6 said, *"It's nice to have these. I feel like we've been trained to drop-down lists, but I did like it, it saves so much time. And I like the description that was given."* Two features, Interactive Legend (F3) and the Annotated Connections Between Text and Visuals such as labels at the top (F2), also received positive feedback from those who used them; however, 4/12 participants did not use them. In interviews, these participants mentioned that they did not expect these elements to be interactive and suggested adding visual cues to encourage viewers to engage with them. In the exit interview, when original WPHPVis participants saw the VisTips version, we received strong positive comments, such as: *"This has exactly what I wanted"*(P5), *"It is a lot more clearer and more comprehensive "*(P19), indicating that VisTips offers advantages and additional clarity. This suggests that benefits perceived by participants stem from embedding explanations directly within the visual representation and interaction flow.

VisTips yielded faster completion times than the original condition in 5/6 tasks, although only Q3 and Q6 reached statistical significance under a Mann-Whitney U test. In Q1, VisTips was faster by an average of 23.4 s, but the difference was not significant (95% CI: [–55.7 s, 8.9 s]; $p = 0.132$). Q2 showed a 7.4 s difference (95% CI: [–18.0 s, 3.2 s]; $p = 0.099$). Q3 was notably quicker with VisTips, by 47.0 s (95% CI: [–78.2 s, –15.8 s]; $p = 0.003$). For Q4, the difference was 16.3 s, not significant (95% CI: [–34.2 s, 1.6 s]; $p = 0.072$). Finally, Q5 showed no difference (95% CI: [–44.8 s, 21.2 s]; $p = 0.224$), and Q6 had a significant 36.3s faster completion time with VisTips (95% CI: [–62.4 s, –10.2 s]; $p = 0.032$).

6.1. Questionnaires

Participants rated VisTips better than the original visualization on all six 5-point statements (with lower scores indicating more favorable responses). Ranked from largest to smallest differences (original minus VisTips), "The visualization supported the tasks effectively" (Q3) was rated 1.78 points higher by participants in the original condition (95% CI: [1.12, 2.44], $p < 0.001$), followed by "It was easy for me to use" (Q1) with a difference of 1.47 points (95% CI: [0.80, 2.14], $p = 0.002$), "I would reuse this visualization design if I had similar data" (Q6) with a difference of 1.22 points (95% CI: [0.55, 1.89], $p = 0.005$), "I enjoyed using the visualization" (Q4) with a difference of 1.17 points (95% CI: [0.51, 1.83], $p = 0.007$), "The information presented in the visualization was clear" (Q2) with a difference of 1.03 points (95% CI: [0.40, 1.66], $p = 0.015$), and "I am confident in the correctness of answers I provided" (Q5) with a difference of 0.74 points (95% CI: [0.12, 1.36], $p = 0.042$). These findings indicate that VisTips was generally rated more favorably, as evidenced by consistently lower scores relative to the original interface. For completeness, a figure comparing participants' ratings across conditions is included in the Supplementary Material.

6.2. Insights into Viewers' Spontaneous Interaction

Spontaneous Interaction as a Stable Sensemaking Strategy: All participants, across both conditions, engaged in spontaneous, unstructured interactions, such as hovering and clicking, particularly during their initial encounters with the visualization, supporting behavioral foundations of our proposed approach. P23 who used the original version said, *"I'm just looking through the options to get a better idea [...]. So, I'm just clicking this up to see what will show up."* Similarly, P20, who used VisTips, said, *I'm playing around to see what each line represents*. Importantly, participants continued to engage in spontaneous interaction across all phases of the sensemaking process, from initial orientation to task completion.

Reshaping Sensemaking: VisTips, as a responsive interface, prompted more interaction as expected [Bea07]. As participants discovered that many elements were interactive, they began actively testing more of them to see what would respond. *"I think after realizing almost everything is interactive, I started hovering more on things just to see like, if I can hover another one and like where it takes me."* (P13) This also meant a shift in how viewers approached sensemaking, as they used each immediate response to construct and refine their understanding of the visualization. Compared to less interactive visualizations, this behavior reflects a more dynamic and feedback-oriented process of understanding.

Facilitating Insight: VisTips facilitated insights as participants were decoding the visualization. In contrast, the original version did not facilitate insight or discovery as effectively: either participants did not mention insights or discovery about the data or it took a longer time in comparison VisTips.

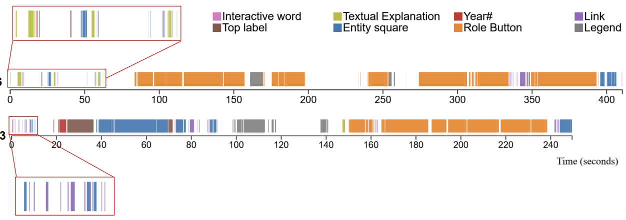

Figure 5: Interaction timelines for two VisTips participants, showing the sequence and duration of their interactions. It shows different paths of sensemaking and how viewers first trying out different elements followed by deep exploration of them.

Temporal Phases of Spontaneous Interaction: In both conditions, participants began sensemaking by randomly hovering over (or clicking in Original Version) various visualization elements. The first phase of spontaneous interaction was characterized by quick and scattered movements, where participants tested different elements to see which ones responded. In the VisTips condition, participants' engagement was not limited to detecting interactivity alone; they also closely observed the changes triggered by

VisTips, such as the appearance of tooltips, highlights, and filtering effects. After this initial round of exploration, participants returned to a subset of elements, this time engaging with them more intentionally. In this second phase, they spent more time on specific elements, carefully inspecting the explanations provided and analyzing how interactions influenced the visualization. These two phases are illustrated in Figure 5, which shows the timelines of interactions for two VisTips participants. This behavior suggests that spontaneous interaction is a layered process: first, viewers probe the interface to gather initial information for understanding how the visualization works and what it represents, and then they refine their comprehension through more deliberate interactions with elements they find relevant or informative.

Sensemaking Paths: Despite the consistency in its occurrence, among all 24 participants, the form that spontaneous interaction took varied in how they moved their cursors, which elements they hovered over or clicked, how long they explored certain features, and in what order they interacted with elements. We observed distinct interaction patterns with no exact repetition, underscoring the inherently personal nature of this behavior (Figure 5). This variability aligns with prior studies which similarly report no fixed global sequence of analytic processes. For example, Isenberg et al. [ITC08] observed substantial temporal variation across participants in studying collaborative visual analysis. Our observations indicate that the sequence in which elements are explored can influence both comprehension and overall experience. P6, who initially interacted with the person bars, later noted upon seeing the title bars and their connections, *"If I started with the Titles first, I would understand it sooner."* This insight suggests that the initial choice of interaction points can set the stage for how effectively users decode the visualization.

Support for Verifying Assumptions: Participants often used VisTips to confirm their initial assumptions, echoing prior findings that viewers seek to validate their interpretations [LKH*15,RTC24, MMF19]. P6, for example, initially attempted to connect people, titles, and firms before realizing that the relationships were structured differently and said: *"I think previously I was trying to connect the person and the title and then the firms that they have published. But firm is related to title individually, and the person is also related to title individually."* Similarly, P13 adjusted their understanding through repeated interaction, saying: *"Oh, so these are women. Oh, I got it now. So these are women that are authors. These are women that are publishers."* In both cases, participants initially formed mental models based on their observations and then verified them by re-engaging with the interactive explanations. This was also observed in the original version of the visualization; however, without VisTips, participants lacked an immediate way to confirm their interpretations. When participants of the original version were shown the VisTips version in the exit interview, they tried to verify their guesses. For example, P19, after interacting with the interactive legend part said, *"oh, so length of each bar does show quantity. This is really helpful because in the previous one [the original Version] It wasn't really clear why they made different lengths."*

6.3. Insights into Design

Direction of Element Expansion: Progressive exploration of aggregated elements (F1) stood out as the most helpful feature of VisTips. In this visualization, each blue bar represents an aggregation of multiple entities, with the bar's length corresponding to the number of entities. When a user hovers over a blue bar in 'person' column, it expands to reveal individual squares, each represnts a person. Moreover, each square further expands vertically to display lines that indicate the connections between that person and the titles to which they contributed (Figure 2a). All VisTips participants recognized the correct interpretation as soon as they saw the horizontal expansion while 3 participants misinterpreted the vertical expansion and assumed that the arrangement, which spans multiple years, indicated that the person was actively publishing throughout that period. In fact the vertical expansion was unrelated to the time scale. For instance, P3, said: *"I think this [the expanded square] is the timespan this author was active, or the time it took to write something because it is covering to these years"*.

Layers of Explanation: For explanatory bubbles provided as part of F1, some participants expressed a desire for additional layers of explanation within the visualization. In particular, they indicated that hovering over the explanatory bubbles, and even over individual words within these bubble (e.g Figure 2a)s, could trigger the highlighting of related visual elements or display further details about the specific term. This behavior suggests that users may require more granular explanatory support to effectively decode the visualization. Moreover, several participants clicked on links expecting the interface to update and reveal more information about the corresponding author or title.

Tutorial or VisTips? The intention behind VisTips was to give viewers the freedom to start with any element and receive explanations as needed. This approach successfully provided flexibility for viewers to build their own sensemaking path by allowing them to interact with elements in any order (shown in Figure 5). Most participants appreciated the explanations being available on interaction, as it allowed them to access information only when they wanted or needed it. For instance, P13 said: *"I think it's great that each line, when you're hovering over them, at least you get some information right then and there."* However, two participants suggested that incorporating more structured guidance could further enhance the experience. For example, P7 said, *"I was a little bit confused when I first entered this interface because I didn't know where should I start."* Also, P3, P7 and p17 asked for visual cues to show the things that would provide explanation on hover. These observations raise this research question: *how can a visualization system balance on-demand explanations with lightweight guidance mechanisms that provide direction?*

Differences in Responses to Overlaid Visual Changes: Participants exhibited varied responses to the visual changes triggered by VisTips, highlighting differences in how users noticed and reacted to the overlaid visual and textual feedback. Eye-tracking data revealed that while 10 VisTips participants immediately noticed highlights and tooltips appearing on hover, two participants either overlooked them or took time to become aware of them. This delay suggests that subtle visual cues may not always be sufficient, and viewers may require additional reinforcement to recognize the provided explanations. One participant even misinterpreted explanations on the links as an error message, as certain words were highlighted in red (shown in Figure 2a). P8 said: *"I thought that is an error, so I ignored it!"* These findings indicate that while immediate feedback supports sensemaking for many users, its effectiveness varies based on individual attention patterns, expectations, and prior experiences with interactive visualizations. This raises a research question: *how* *can visual and textual feedback be designed to accommodate diverse attention styles and avoid misinterpretation?*

7. Discussion

We introduced an interaction paradigm, SEA, as a complementary approach to existing onboarding strategies, such as tutorials. Its novelty lies in empowering viewers to build their own learning paths by treating spontaneous interactions as moments to deliver lightweight, transient explanations through visual and interaction mechanisms (e.g. visual transformation.)

The constructs we introduced to describe SEA outline an initial design space, but they do not exhaust its possibilities. Future work could explore a broader range of triggers including gaze, scroll, or touch as well as alternative explanation modalities such as animation or audio. Additional dimensions may also be worth exploring; for example, incorporating viewer *intent*, or what a viewer seeks to accomplish through specific interactions. Exploration into using the approach with different data (e.g. spatial data) or on different devices (e.g., mobile, AR/VR) may further expand the design space. As such, SEA is not a fixed template, but rather a flexible design lens that invites further experimentation across contexts. Having said that, our findings reveal design considerations that should be taken into account when applying SEA in practice.

Designing for unstructured exploration with awareness of visual flow. We deliberately avoided directing viewers' attention with eye-catching elements. Instead, we designed the visualization to maintain a balanced aesthetic, allowing users to naturally work with visualization elements. This decision reflects our intention to empower users to build their own paths of understanding. However, it is important to recognize that natural tendencies often shape how interactions unfold. For example, viewers from Western countries often naturally start exploring from left to right and top to bottom [SH05]. Therefore, being mindful of these natural tendencies is crucial to ensure that the order in which elements are interacted with has minimal impact on the sensemaking process.

Explanation design can be nuanced and complex. These findings highlight the importance of carefully designing not only what explanations are shown, but when and how they are revealed, as they might add incorrect information or confuse viewers. For example, in VisTips, the order in which cumulative elements are expanded may influence interpretations (Direction of element Expansion). This raises design questions about how to align dynamic explanations, especially visual transformation, with expectations to minimize potential misinterpretations. Additionally, our observations about viewers wanting explanations (Layers of Explanation) on the explanations, align with the "overview first, zoom and filter, details on demand" principle [Shn03], suggesting an opportunity to incorporate multi-layered explanations that unfold progressively. Such interactions could support both early exploration and deeper sensemaking as viewers gain familiarity with the data. Future research could examine how the sequencing and layering of explanations influence comprehension, to what extent these choices shape the paths users take through complex visualizations, and whether personalization of the explanation responses is helpful.

7.1. Transferability of the Paradigm: Lessons Learnt

While our implementation of SEA through VisTips was specific to a complex network visualization (WPHPVis), we believe the conceptual design principles behind Serendipitous Explanations offer a flexible lens for thinking about interaction-driven comprehension across diverse visualization contexts. In particular, visualizations that present unfamiliar encodings, or that have steep learning curves can benefit from treating spontaneous viewer interactions as entry points for embedded explanation.

For SEA to be applicable, the visualization must expose meaningful opportunities for interaction without conflicting with existing controls. SEA could be implemented as a layer in visualization authoring environments (e.g. Observable, D3.js), where interaction events are already captured and can be repurposed to trigger integrated in-situ explanations. In contrast, more dynamic features, particularly transforming cumulative elements to reveal individual components, present notable challenges. Yet despite their complexity, these transformations proved to be the most effective feature of VisTips (Figure 3), offering immediate and intuitive comprehension by visually externalizing the structure behind the encoding.

Since the paradigm relies on spontaneous behaviors like hovering and clicking as explanation triggers, these interactions must be available and not already occupied by core visualization functionality (e.g., tooltips used for data values, clicks used for filtering). Furthermore, prior studies [BM19, IH00, GGW13] have shown that users naturally experiment with interactions such as scrolling, panning, and zooming, when exploring complex visual representations, suggesting that it might be valuable for VisTips to support these behaviors as well. However, this poses a new design challenge, as interactions such as zooming or panning may conflict with hover-triggered or click-based explanations. For example, zooming into an element to examine its finer details might also trigger a hover explanation, creating a conflict between navigation and explanation. Addressing these challenges requires developing context-aware interaction strategies that can seamlessly integrate explanatory feedback without interfering with primary analytic tasks.

7.2. Limitations

VisTips, as an instance of the proposed approach, contains many variations of serendipitous explanations, but by no means exhausts the possibilities. While the success of our initial version of VisTips is promising, we recognize that other instances can be designed to suit different visualizations, situations and needs. Also, our study only included 24 participants (16 students). While this sample size may limit the statistical power, it was sufficient to observe qualitative themes and patterns and to reach thematic saturation.

8. Conclusions

To complement existing structured onboarding approaches, we introduced the Serendipitous Explanations Approach (SEA), which frames spontaneous viewer interactions, such as hovering on elements, as opportunities to deliver explanations within visualizations. Through VisTips, an instantiation of SEA, we advanced understanding of spontaneous interaction, and demonstrated how SEA supports and influences visualization sensemaking. Visual explanations such as transformations or highlights, were especially effective in facilitating comprehension. We invite the visualization community to help expand this approach by exploring new triggers, explanation techniques, and interaction modalities, to collectively build a rich design space for Serendipitous Explanations.

Acknowledgment: We acknowledge and thank the Natural Sciences and Engineering Research Council of Canada (NSERC) [RGPIN-201907192], and the Canada Research Chair in Data Visualization [CRC2019-00368].

References

[BE14] BLASCHECK T., ERTL T.: Towards analyzing eye tracking data for evaluating interactive visualization systems. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization* (2014), pp. 70–77. 3

[Bea07] BEALE R.: Supporting serendipity: Using ambient intelligence to augment user exploration for data mining and web browsing. *International Journal of Human-Computer Studies 65*, 5 (2007), 421–433. 8

[BM19] BUCHANAN G., MCKAY D.: One way or another i'm gonna find ya: The influence of input mechanism on scrolling in complex digital collections. In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)* (2019), IEEE, pp. 287–296. 10

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D^3 data-driven documents. *IEEE transactions on visualization and computer graphics 17*, 12 (2011), 2301–2309. 6

[BOZ*14] BROWN E. T., OTTLEY A., ZHAO H., LIN Q., SOUVENIR R., ENDERT A., CHANG R.: Finding waldo: Learning about users from their interactions. *IEEE Transactions on visualization and computer graphics 20*, 12 (2014), 1663–1672. 3

[BS99] BRIDGE P. D., SAWILOWSKY S. S.: Increasing physicians' awareness of the impact of statistics on research outcomes: comparative power of the t-test and wilcoxon rank-sum test in small samples applied research. *Journal of clinical epidemiology 52*, 3 (1999), 229–235. 7

[CS90] CORBIN J. M., STRAUSS A.: Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology 13*, 1 (1990), 3–21. 7

[CYC*23] CHUNDURY P., YALÇIN M. A., CRABTREE J., MAHURKAR A., SHULMAN L. M., ELMQVIST N.: Contextual in situ help for visual data interfaces. *Information Visualization 22*, 1 (2023), 69–84. 3

[DHF*24] DHANOA V., HINTERREITER A., FEDIUK V., ELMQVIST N., GRÖLLER E., STREIT M.: D-tour: Semi-automatic generation of interactive guided tours for visualization dashboard onboarding. *IEEE Transactions on Visualization and Computer Graphics* (2024). 2

[dMvO96] DE MUL S., VAN OOSTENDORP H.: Learning user interfaces by exploration. *Acta psychologica 91*, 3 (1996), 325–344. 3

[DRRD12] DÖRK M., RICHE N. H., RAMOS G., DUMAIS S.: Pivotpaths: Strolling through faceted information spaces. *IEEE transactions on visualization and computer graphics 18*, 12 (2012), 2709–2718. 3

[DWH*22] DHANOA V., WALCHSHOFER C., HINTERREITER A., STITZ H., GROELLER E., STREIT M.: A process model for dashboard onboarding. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 501–513. 1

[Fie13] FIELD A.: *Discovering statistics using IBM SPSS statistics, 4th edition*. Sage publications limited, 2013. 6

[GGW13] GLUECK M., GROSSMAN T., WIGDOR D.: A model of navigation for very large data views. In *Proceedings of Graphics Interface 2013*. 2013, pp. 9–16. 10

[GWL*19] GÄBLER J., WINKLER C., LENGYEL N., AIGNER W., STOIBER C., WALLNER G., KRIGLSTEIN S.: Diagram safari: A visualization literacy game for young children. In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts* (2019), pp. 389–396. 2

[Hei91] HEIN G. E.: Constructivist learning theory. *Institute for Inquiry 14* (1991). 2, 3

[HS25] HOQUE N., SULTANUM N.: Dashguide: Authoring interactive dashboard tours for guiding dashboard users. In *Computer graphics forum* (2025), vol. 44, Wiley Online Library, p. e70107. 1

[IH00] IGARASHI T., HINCKLEY K.: Speed-dependent automatic zooming for browsing large documents. In *Proceedings of the 13th annual ACM symposium on User interface software and technology* (2000), pp. 139–148. 10

[ITC08] ISENBERG P., TANG A., CARPENDALE S.: An exploratory study of visual information analysis. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2008), pp. 1217–1226. 9

[KL16] KWON B. C., LEE B.: A comparative evaluation on online learning approaches using parallel coordinate visualization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), pp. 993–997. 2

[LFH17] LAZAR J., FENG J. H., HOCHHEISER H.: *Research methods in human-computer interaction*. Morgan Kaufmann, 2017. 6

[LKH*15] LEE S., KIM S.-H., HUNG Y.-H., LAM H., KANG Y.-A., YI J. S.: How do people make sense of unfamiliar visualizations?: A grounded model of novice's information visualization sensemaking. *IEEE transactions on visualization and computer graphics 22*, 1 (2015), 499–508. 1, 2, 3, 4, 9

[MAC*20] MAHMUD S., ALVINA J., CHILANA P. K., BUNT A., MCGRENERE J.: Learning through exploration: how children, adults, and older adults interact with a new feature-rich application. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–14. 3

[Mar06] MARCHIONINI G.: Exploratory search: from finding to understanding. *Communications of the ACM 49*, 4 (2006), 41–46. 3

[MM25] MACKAY W. E., MCGRENERE J.: Comparative structured observation. *ACM Trans. Comput.-Hum. Interact. 32*, 2 (Apr. 2025). URL: `https://doi.org/10.1145/3711838`, `doi:10.1145/3711838`. 7

[MMF19] MA J., MA K.-L., FRAZIER J.: Decoding a complex visualization in a science museum–an empirical study. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 472–481. 1, 2, 4, 9

[MTW*12] MOERE A. V., TOMITSCH M., WIMMER C., CHRISTOPH B., GRECHENIG T.: Evaluating the effect of style in information visualization. *IEEE transactions on visualization and computer graphics 18*, 12 (2012), 2739–2748. 3

[Mun14] MUNZNER T.: *Visualization analysis and design*. CRC press, 2014. 5

[MVFM22] MASSON D., VERMEULEN J., FITZMAURICE G., MATEJKA J.: Supercharging trial-and-error for learning complex software applications. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022), pp. 1–13. 3

[PWM10] POHL M., WILTNER S., MIKSCH S.: Exploring information visualization: describing different interaction patterns. In *Proceedings of the 3rd BELIV'10 Workshop: Beyond time and errors: novel evaluation methods for information visualization* (2010), pp. 16–23. 3

[RJPL16] REDA K., JOHNSON A. E., PAPKA M. E., LEIGH J.: Modeling and evaluating user behavior in exploratory visual analysis. *Information Visualization 15*, 4 (2016), 325–339. 3

[RTC24] REZAIE M., TORY M., CARPENDALE S.: Struggles and strategies in understanding information visualizations. *IEEE Transactions on Visualization and Computer Graphics 30*, 6 (2024), 3035–3048. `doi:10.1109/TVCG.2024.3388560`. 1, 2, 3, 4, 9

[SB19] SCHNEIDER K., BERTOLLI L. M.: Video variants for crowdre: How to create linear videos, vision videos, and interactive videos. In *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)* (2019), pp. 186–192. `doi:10.1109/REW.2019.00039`. 2

[SGA22] STOIBER C., GRASSINGER F., AIGNER W.: Abstract and concrete materials: What to use for visualization onboarding for a treemap visualization? In *Proceedings of the 15th International Symposium on Visual Information Communication and Interaction* (New York, NY, USA, 2022), VINCI '22, Association for Computing Machinery. URL: `https://doi.org/10.1145/3554944.3554949`, `doi:10.1145/3554944.3554949`. 2

[SGP*19] STOIBER C., GRASSINGER F., POHL M., STITZ H., STREIT M., AIGNER W.: Visualization onboarding: Learning how to read and use visualizations. 2

[SH05] SPALEK T. M., HAMMAD S.: The left-to-right bias in inhibition of return is due to the direction of reading. *Psychological Science 16*, 1 (2005), 15–18. 10

[Shn03] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*. Elsevier, 2003, pp. 364–371. 10

[SWG*21] STOIBER C., WALCHSHOFER C., GRASSINGER F., STITZ H., STREIT M., AIGNER W.: Design and comparative evaluation of visualization onboarding methods. In *Proceedings of the 14th International Symposium on Visual Information Communication and Interaction* (New York, NY, USA, 2021), VINCI '21, Association for Computing Machinery. URL: https://doi.org/10.1145/3481549.3481558, doi:10.1145/3481549.3481558. 2

[SWP*22] STOIBER C., WALCHSHOFER C., POHL M., POTZMANN B., GRASSINGER F., STITZ H., STREIT M., AIGNER W.: Comparative evaluations of visualization onboarding methods. *Visual Informatics 6*, 4 (2022), 34–50. 1, 2

[TAM17] TARAMIGKOU M., APOSTOLOU D., MENTZAS G.: Supporting creativity through the interactive exploratory search paradigm. *International Journal of Human–Computer Interaction 33*, 2 (2017), 94–114. doi:10.1080/10447318.2016.1220104. 3

[TGF13] TOGELIUS J., GUSTAFSSON FRIBERGER M.: Bar chart ball, a data game. In *Foundations of Digital Games (FDG), Chania, Crete, Greece (2013)* (2013), Society for the Advancement of the Science of Digital Games (SASDG), pp. 451–452. 2

[THC12] THUDT A., HINRICHS U., CARPENDALE S.: The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the SIGCHI conference on human factors in computing systems* (2012), pp. 1461–1470. 2, 3

[TRL*24] TAGHIPOUR P., REZAIE M., LEVY M., SHERMER T., CARPENDALE S.: Supporting exploration of women's print history project data via interactively constructing networks of interest. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces* (New York, NY, USA, 2024), AVI '24, Association for Computing Machinery. URL: https://doi.org/10.1145/3656650.3656697, doi:10.1145/3656650.3656697. 5, 6

[VG22] VEARS D. F., GILLAM L.: Inductive content analysis: A guide for beginning qualitative researchers. *Focus on Health Professional Education: A Multi-Professional Journal 23*, 1 (2022), 111–127. 7

[WCS*22] WANG J., CAI X., SU J., LIAO Y., WU Y.: What makes a scatterplot hard to comprehend: data size and pattern salience matter. *Journal of Visualization* (2022), 1–17. 1, 2

[Wom] Women's print history project. https://womensprinthistoryproject.com/. Accessed: 2025-03-1. 5

[YEB16] YALÇIN M. A., ELMQVIST N., BEDERSON B. B.: Cognitive stages in visual data exploration. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), pp. 86–95. 3