NetworkNarratives: Data Tours for Visual Network Exploration and Analysis

Wenchao Li
The Hong Kong University of Science
and Technology
Hong Kong SAR, China
wlibs@connect.ust.hk

Sarah Schöttler
University of Edinburgh
Edinburgh, United Kingdom
sarah.schoettler@ed.ac.uk

James Scott-Brown
University of Edinburgh
Edinburgh, United Kingdom
james@jamesscottbrown.com

Yun Wang
Microsoft Research Asia
Beijing, China
wangyun@microsoft.com

Siming Chen
Fudan University
Shanghai, China
siming.chen@fudan.edu.cn

Huamin Qu
The Hong Kong University of Science
and Technology
Hong Kong SAR, China
huamin@cse.ust.hk

Benjamin Bach
University of Edinburgh
Edinburgh, United Kingdom
bbach@inf.ed.ac.uk

ABSTRACT

This paper introduces semi-automatic data tours to aid the exploration of complex networks. Exploring networks requires significant effort and expertise and can be time-consuming and challenging. Distinct from guidance and recommender systems for visual analytics, we provide a set of *goal-oriented* tours for network overview, ego-network analysis, community exploration, and other tasks. Based on interviews with five network analysts, we developed a user interface (NetworkNarratives) and 10 example tours. The interface allows analysts to navigate an interactive slideshow featuring facts about the network using visualizations and textual annotations. On each slide, an analyst can freely explore the network and specify nodes, links, or subgraphs as seed elements for follow-up tours. Two studies, comprising eight expert and 14 novice analysts, show that data tours reduce exploration effort, support learning about network exploration, and can aid the dissemination of analysis results. NetworkNarratives is available online, together with detailed illustrations for each tour.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**.

KEYWORDS

Guided exploration, network visualization

1 INTRODUCTION

Large, dense, and multivariate (potentially including node types or geographic locations, and link types, weights, times, or directions) relational datasets (networks) pose challenges to exploratory data analysis. Visualization interfaces for interrogating and exploring networks have grown increasingly sophisticated in order to support the richness of potential questions [40]. Tools such as Gephi [11], Palladio [37], TempoVis [3], or Visone [12], provide different visual encodings, interaction, analysis metrics, and multiple visualization types, such as adjacency matrices, timelines, and maps that are sometimes presented as multiple coordinated views [8].

Although powerful when used by an experienced analyst, feature-rich user interfaces present challenges for novice analysts who are required to learn possible interactions, understand the aim, perform interactions, and keep track of everything. Furthermore, significant time can be taken up by repeating steps, applying them to different datasets, keeping track of one's exploration, employing layout and exploration strategies, and undoing interactions in case of mistakes. In such free-form exploration interfaces, analysts can become lost or overwhelmed [84] or make analysis errors, such as succumbing to the *drill-down fallacy* [41]. For a novice analyst, a rich set of tool features can result in a steep learning curve that requires cognitive effort to understand each feature and its affordances and effects [16]. As highlighted in a recent study [4], the open-ended nature of exploration can be overwhelming to novice analysts who may not know what information can be gleaned from

a network and what questions can be answered with network visualization. Therefore, creating effective exploration strategies and learning about network exploration without appropriate training or experience are challenging.

In this work, we explore the idea of semi-automatic data tours to aid in network exploration. The idea of data tours goes back to Asimov's Grand Tour for multivariate data [7] and has been recently described in a theoretic framework [46]. We obtain further inspiration from ideas on guidance (e.g., as implemented in tools such as SocialAction [57] and Small Multiples [73]), recommender systems [36], and data-driven storytelling (notably graph comics [9] and interactive slideshows [63]). In our case, a *data tour* walks an analyst through their network, similar to viewing a slideshow presentation created by another analyst (Figure 1). Each slide in a tour (e.g., panels above the red line representing the *Network Overview* tour) shows a specific piece of information about the network as explained in the caption. Our data tours aim to lower the barrier for novice analysts to learn interaction and exploration strategies and provide quick overviews of unknown datasets to expert analysts. Our data tours are defined by three main characteristics:

First, unlike existing recommender systems (e.g., [65, 75]), our approach is *goal-driven*, i.e., a data tour in our case is best thought of as a template representing an exploration strategy or story. We loosely describe an exploration or analysis strategy as *"a set of information in a purposeful order, selected with the goal of providing insights into a dataset."* For example, an analyst selects a subgraph and learns about the number of its nodes and density, as well as the most connected nodes and relations to the rest of the network (Figure 1). To the best of our knowledge, this paper is the first to describe a practical approach to data tours in the domain of network analysis.

Second, tours are primarily *sequential*, allowing an analyst to easily flip through facts at their own pace while reducing the cognitive load imposed by decision-making about navigation and filtering. At any point, an analyst can take control, freely explore the network, and embark on a new tour or detours by inserting facts from NetworkNarratives' recommender engine (Figure 1(D)).

Finally, we implement 10 complementary example data tours that are *linked* so an analyst can pivot between them to change the focus of his exploration. For example, when a slide focuses on a specific node, the user is offered the opportunity to start a tour on the ego-network exploration of that node (Figure 1(B)+(C)). Although not all presented facts in our data tours might be of interest to an analyst, we agree with Tukey regarding the importance of *"notice what we never expected to see"* [72, Preface], *"[to find] nothing, is a definite step forward"* [72, Preface], and to *"give users a chance to think of initial questions [to help] them get started"* [52, 71].

Our concept of data tours is defined with **six design goals** in mind (subsection 3.2). Our 10 individual tours are designed in collaboration with five network analysis experts with backgrounds in social science, history, epidemiology, and archaeology. We have implemented these tours in our user interface and recommender system *NetworkNarratives*, which allows users to choose tours, navigate tours, and freely explore the network. Qualitative feedback is obtained from eight network analysis experts who explored their own data with NetworkNarratives. In addition, a comparative study

with 14 novice analysts that suggests that our tours save time during the exploration process and that a goal-driven approach can make the exploration more accessible (section 7): tours provide a simple set of entry points and allow analysts to choose from a well-defined set of tours, each representing a specific analysis goal.

In summary, our contributions are as follows:

(1) the concept of data-driven data tours for network analysis;
(2) 10 extensible fully-implemented data tours (section 4), including 102 individual facts (subsection 3.3) for multivariate, temporal, and geographic networks;
(3) *NetworkNarratives*, an interactive user interface to experience data tours (section 5), which is publicly available and can be used as either a standalone application or an extension to *The Vistorian* [4] (documentation and video demos are also available online: https://networknarratives.github.io);
(4) two studies with 8 network analysis experts and 14 novice analysts respectively, which evaluate the usefulness and future potential of data tours and the NetworkNarratives system (section 7).

2 RELATED WORK

2.1 Data Tours

Data Tours were first introduced by Asimov, whose Grand Tour [7] was a computer-generated overview of a multivariate dataset that transitioned between scatterplots showing different projections [72]. Such tours have been created for time-varying data [85] and multivariate data [28]. Two recent approaches model individual facts about a dataset as nodes of a graph, with links between graphs as possible relations. In Graphscape [36], these nodes are connected to form tours (or stories). Similarly, Mehta et al. [46] described a theoretical framework for a hierarchical structure of facts for data tours, including staged transitions between individual facts/nodes.

In this work, we explore data tours for the exploration and analysis of networks. Closest to our data tours is the Systematic Yet Flexible (SYF) concept by Perer and Shneiderman [55, 56]. This approach combines the structured exploration of social networks through a predefined set of seven *steps* (overview, rank nodes, rank edges, plot nodes, plot edges, find communities, edge types), each featuring a set of *tasks*, such as different node degree rankings. Each step is supported by a visualization such as a node-link diagram or a ranking visualization. NetworkNarratives builds on this concept by diversifying the notion of a single data tour to multiple (10) tours, each focusing on a specific goal such as exploring an ego-network, comparing two subgraphs, or exploring temporal and geographic networks. In addition, NetworkNarratives tries to minimize manual interaction and provides detours on demand.

2.2 Onboarding, Guidance, and Storytelling

Onboarding and guidance are closely related concepts that aim to help users of (visualization) systems, tools, and techniques. Onboarding can include step-by-step wizards, guided tours, video-based tutorials, help centers or overlays [68, 69], or cheatsheets [78]. Here, data tours can be seen as a form of onboarding. Instead of onboarding onto a specific tool, system or visualization technique, our
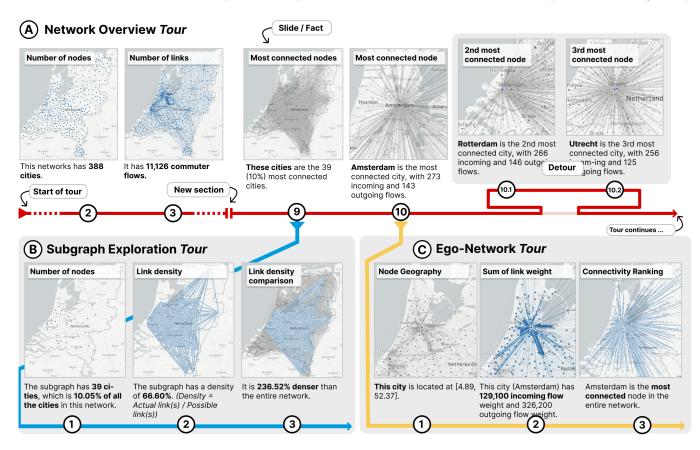
Figure 1: Conceptual illustration of data tours in NetworkNarratives. An initial data tour of Network Overview (A, red line) showing facts (on slides) about the network. Each slide has a title and textual description. Numbers in circles indicate the number of facts in the tour. At times, e.g., when discussing a node or subgraph, a user can pivot to related tours (B, C) about that specific node or subgraph. Detours (D) include additional slides on demand for any given tour.

system introduces (novice) analysts to questions about networks and methods for visual exploration.

Like guidance systems [18–20, 66], data tours guide a user *while* they are using a system or tool, rather than providing general *a priori* resources and knowledge similar to on-boarding. Systems and techniques for guidance can be categorized based on whether they merely help *orient* a user through overview and aid in building a mental map, suggest options to choose from (*directing*), or *prescribe* views and analyses [18] through storytelling [62]. Data tours prescribe information and views that can be easily followed by clicking through a pre-defined sequence of visualization views. Guidance for the exploration of networks [32, 45, 58] has been implemented by suggesting nodes (*orienting*) that may be of interest to a user based on his current selection [21, 26]. These systems focus on analytical and operational knowledge, i.e., understanding analysis methods and the respective tools and systems. We see our data tours as a first step toward *methodological knowledge* for exploration and analysis to explain questions, goals, and specific steps, especially when the user lacks any goals or hypotheses to inform their actions [4].

Provenance systems are similar to *guidance* systems in that both represent many different possible visualizations of a dataset. However, guidance systems deal with visual representations that a user could potentially look at in the future, and provenance systems (e.g., StoryFacets [53]) are concerned with the visualizations that have already been viewed. Data tours could be created from a previous exploration and analysis history by making insights generic, e.g., by turning them into templates and re-applying them to different datasets.

Data tours also draw inspiration from data-driven storytelling, specifically, graph comics [9] and slideshows [63]. We follow the pattern of an interactive slideshow [63] that combines elements of author-driven storytelling (the slides and tours) and reader-driven storytelling (interactive exploration, selecting nodes and graph for further tours, and navigating tours). Data tours could be seen as paths through a network, and analysts are encouraged to create their own tours. However, we did not design NetworkNarratives as a story authoring tool nor an automatic storytelling system. Actual data-driven storytelling requires information about the context and audience to be effective.

2.3 Data Fact Recommendation

Automatic fact extraction has been used in guidance to review large datasets [39], start the exploration loop [80], support communication (e.g., [75, 76, 85]), and avoid drill-down fallacies [41]. The problem of automatically visualizing facts from data can be subdivided into: (i) identifying individual data facts (often referred to as *insights*); (ii) suggesting visualizations for each fact; and (iii) presenting facts and visualizations alongside each other. NetworkNarratives implements some of these concepts.

First, systems to extract data facts have been summarized by Law et al. [39]. Data facts can be selected using statistical techniques through a separate procedure to generate facts for each of a predetermined list of fact types. Facts can then be ordered based on a ranking score that is intended to represent subjective relevance, importance, or degree of interest (e.g., [27, 65, 75]). To the best of our knowledge, all these systems explicitly focus on tabular data, rather than network or relational data as explored in this work. The facts we consider include common network metrics as well as topological features.

Second, a number of systems automatically suggest either a single visualization or a set of visualizations that most appropriately depict a dataset (e.g., [30, 42, 48, 73, 86]); these systems have been summarized elsewhere [39]. Visual encodings can be chosen based on either rules expressing guidelines using constraint satisfaction [48] or machine learning methods trained on examples of human choices [30, 86]. The systems allow users to interactively browse chart recommendations [81] or receive them as notifications [23]. In NetworkNarratives, we initially limit ourselves to node-link diagrams, rather than automatically choosing from a range of visualization types. Other visualization techniques, such as adjacency matrices for dense graphs, are likely to require additional explanations (e.g., [44, 67, 78]) to enable new users to read them properly and obtain *visual insights* [39].

Finally, facts and visualizations must be related and presented to the user, possibly in a form that enables interactive exploration and personalization. For example, researchers have designed interfaces that group insights into panels based on category [24, 27], display dashboards [34], or generate infographic-like fact sheets of visualizations featuring textual explanations [75]. To organize data facts into sequences, graph-based approaches create a similarity graph from all facts and consequently select a sequence of visualizations (a path through the graph) based on minimizing edge weight [36, 85]. Data facts can also be ordered using the complex technique of logic-oriented Monte Carlo tree search [65]. An alternative approach is to apply the same type of visualization to all variables in a dataset and then select those that show the most interesting patterns: this is the route taken by scagnostics [51, 79] and GRID [64] for tabular data, and magnostics for adjacency matrices representing network data [14].

All these approaches are predominantly *data-driven*, in that they start by analyzing the data and then ask the user to navigate and express their preferences. By design, these systems have only a loose notion of the inherent human factors underlying exploration, such as research questions and methodologies (whether formal or informal). Our approach is based on sequential high-level *goals* that prescribe a set of facts and their specific sequence presented in

an interactive slideshow [63]. Our templates are human-created to ensure that they target specific network analysis goals and present data facts in a meaningful sequence.

3 DESIGN GOALS AND CHARACTERISTICS

This section describes the methodology of informing data tours and NetworkNarratives (subsection 3.1), followed by the design goals (subsection 3.2) and the resulting main characteristics of our approach (subsection 3.3).

3.1 Research Methodology

Data tours and the NetworkNarratives system were established by a three-step process of literature review, domain expert interviews, and iterative prototyping and design.

1) We examined existing task taxonomies to obtain a systematic overview of facts and information relevant to network exploration [3, 40, 49, 83]. We found 102 facts, and two of the authors formally categorized them by applying tags, such as *centrality*, *temporal*, and *connectivity*. Although we originally aimed for a clear taxonomy of facts, we found that the resulting taxonomy was too ambiguous and not helpful for our work. The tags were ultimately used to search for facts when we created data tours (section 6).

2) We reviewed written reports in papers and blog posts describing network analyses and publications conducted by domain experts in different subject areas (e.g., [1, 38]). From these reports, we extracted metrics and insights and noted the order in which they were reported, yielding an initial collection of data tours to power NetworkNarratives' recommender engine.

3) We interviewed five domain experts in network analysis. The individual interviews lasted one hour on average and focused on exploration and analysis goals, current workflow methods, and tools used. Four of the five analysts had a background in the Humanities and Social Sciences and analyze networks on a daily basis: *Arch* (Associate Prof.) analyzes historical transport networks, focusing on transportation costs (link weight) within the network; *Hist1* (Prof.) explores historical social networks extracted from newspapers and letters; *Health* (Prof.) explores networks of collaborations between researchers and how they influence the policy advice that they provide about disease response; and *Soc1* (Assistant Prof.) and *Soc2* (Assistant Prof.) investigate community structures in the network of interactions between users on social media sites, a task that often involves working with multiple link and node types (both *Soc1* and *Soc2* have a strong quantitative background). Across all interviews, we identified tasks (e.g., understanding paths, exploring ego-networks) and exploration strategies (e.g., following a path along a geographic feature such as a river). Experts confirmed that existing free-form approaches were tedious to use, especially for repetitive steps and routines.

3.2 Design Goals

The design goals were based on our conversations with analysts, the literature, and our own experience in working with network analysts over many case studies.

G1: LEARN—Introduce exploration strategies, goals, and concepts to novice analysts—In our interviews, we saw that exploration is rarely entirely open, and is typically influenced by factors including high-level research goals and prior knowledge about the data, as well as personal analysis protocols and methodologies. For example, *Soc1* and *Soc2* start their analysis by calculating overall network metrics, such as density, fragmentation, average path length, and average degree. They then focus not only on individual nodes, their centrality metrics, and ego-networks, but also on smaller communities within the network. *Hist1* and *Arch* expressed minimal interest in general metrics and instead focus on the detailed exploration of geographic regions (*Arch*) or on individual nodes and their ego-networks, individual links and their attributes, and specific time slices of a dynamic network (*Hist1*).

In an open-ended exploration, novice analysts might lack specific high-level **goals** when exploring a network [4]. They might also be unfamiliar with network **concepts** required to "decode" information from networks (such as clusters and communities, node degrees, the shortest paths, and link weight). To achieve an exploration goal, analysts need **strategies** that they can apply through interactions and reading patterns from visualization (such as selecting nodes, calculating metrics, searching for elements with specific characteristics, and comparing graph elements [40]). Data tours can automate and demonstrate some of these concepts by exemplifying exploration for any given dataset.

G2: REDUCE—Reduce cognitive and manual exploration effort—All the analysts reported using and frequently switching between multiple tools to alternate between network visualizations (for topology and overview tasks) and metric calculation (for analysis tasks). These tools included *igraph* [22] and *tidygraph* [54] for calculating network analysis metrics and producing static visualizations, *Visone* [12] and *Gephi* [11] for interactive visualizations, GIS software for geographic data, and tools that they have written themselves. Switching tools interrupts work and requires piecing together information from different representations. For example, an analyst might use a script to find the node with the highest degree, and then switch to a network visualization tool to explore its context (*"I use visualization for initial exploration, but to actually go into the analytics is difficult because it's messy"* [*Soc1*]). The analysts frequently export subnetworks as separate datasets for further analysis but report difficulties keeping track of these exported files. In addition, an exported subnetwork cannot easily be linked back to its context in the main network: *"[these are] important needs, now that we have these very large databases"* [*Arch*].

Data tours can free an analyst from the majority of manual and cognitive labor required: a) interactions (e.g., selecting a time range, panning and zooming to find an element or visual pattern, or hovering over nodes to show their labels), b) decisions (e.g., what information to look at/what question to ask), c) visual search (e.g., finding the node with the highest degree, or finding the strongest or weakest link), and d) visual counting tasks (e.g., counting the degree of a node, or the number of unconnected nodes).

G3: REPEAT—Repeat routine explorations—Data tours can easily be repeated and reapplied to different networks. This feature can be useful when an analyst wants to repeat an analysis workflow with an updated dataset, apply the same exploration to subsets (e.g., connected components or clusters) of the complete dataset,

or compare multiple versions of the same network (e.g., obtained by applying different filtering nodes and links). One participant reported *"we ended up with tens of thousands of all sorts of networks"* [*Hist1*].

G4: BALANCE—Balance prescription and agency—Data tours need to be inspiring, not restricting. Similar to guidelines, they should give direction and provide detailed steps into that direction. Analysts may change their priorities and shift their high-level goals: *"the more you produce networks, the more you have a chance to really get lost; you don't know exactly what you have"* [*Hist1*].

G5: SURPRISE—Support serendipitous discovery—Despite being goal-oriented, an exploration process needs to remain open to serendipity and unexpected discoveries. Data tours should include a range of information (time, link-weight, and isolated nodes), especially when networks are large and multivariate: *"with my current tool, I use random [strategies] and found it difficult to scale [exploration] at a reasonable size"* [*Soc1*].

G6: TRANSPARENCY—Keep information in tours and tour structure simple and transparent—Our data tours follow regular templates scripted by a human author (section 6). We want to avoid complex recommender models whose decisions might not be transparent to the analyst. Our tours are predictable in that they follow this human script, rather than a user model that tries to learn a user's intentions or can trap an analyst in a "recommender bubble".

Tours for other datasets, audiences, and purposes might be designed with different goals in mind and implement different design decisions.

3.3 Main Characteristics at a Glance

The design goals **G1**–**G6** led us to the following design decisions. Figure 1 displays three example tours and their relationships, each one shown along a colored line: *Network Overview* (red, top), *Subgraph Exploration* (blue, left), and *Ego-Network Analysis* (yellow, right). The Network Overview tour starts by explaining the number of nodes and links in the network, followed by the most connected nodes and the most connected node. In the following section, we explain the main tour concepts illustrated in Figure 1.

Goal-Oriented. Based on expert interviews and our own experience, we designed 10 complementary tours described in detail in section 4. In response to design goals LEARN and REDUCE, each tour has a specific goal that is expressed by its title and abstract. For example, the goal of the *Network Overview* tour (Figure 1(A)) is to provide an overview of the most important facts about this network, such as the number of links and nodes, clusters, or link density. Goals can help inform an individual which tour to choose and why to embark on a specific tour.

Facts and Slides. Each tour is made of a sequence of *facts*. A fact is a piece of information about the network such as the number of nodes and links, network density. Facts are expressed as templates with the respective values calculated from the dataset (e.g., `This network has N nodes`). In NetworkNarratives, a fact is visually represented as a *slide* (panels in Figure 1) that shows a node-link visualization of that network, potentially highlighting relevant nodes and links, and describing the fact in a caption.

A caption also explains unknown concepts, e.g., explaining *link-density* alongside a link to a related web resource (LEARN). The number and sequence of facts in each tour are predetermined to serve a goal (LEARN). Depending on the data characteristics (e.g., temporal, geographic, link weights), irrelevant slides are removed from the tour by identifying the specific tags of topics covered in a fact. Once a tour has been developed as a template (section 6), it can be applied to many different datasets (REPEAT). Facts are automatically populated with metrics and visualizations computed from the provided data.

Sequence and Structure. To facilitate understanding of the tour and navigation (TRANSPARENCY), slides (facts) for each tour are presented in a single, sequential order (Figure 1). Slideshows are designed to provide cohesion among the facts and reflect the specific steps of an exploration strategy (LEARN). For example, a data tour might start with a high-level overview of basic network metrics and then explain one aspect in more detail while potentially linking to other tours. Alternatively, it might start with a single node and then zoom out to explore its ego-network and further to progressively introduce the surrounding nodes. A data tour could also follow a set of nodes on a path, highlighting connections between them (section 4). Facts about a similar theme are grouped into *sections* to structure a tour into meaningful units and support fast navigation between sections (REDUCE). For example, the Network Overview tour (Figure 1) is subdivided into the sections *Overview* (e.g., number of nodes and links, network density, and most connected nodes), *Link information* (e.g., strongest link and weakest link), *Nodes and centralities*, and *Node clusters*.

Navigation. Navigation through the facts and sections of a tour is mainly linear (REDUCE). At specific points, an analyst can pivot into a related tour. For example, the 4th slide in the red tour in Figure 1 mentions the most connected node in that network. An analyst can now decide to launch a tour on this node's ego-network, that is, launching the *Ego-Network Analysis* tour (yellow). Meanwhile, the focus of NetworkNarratives is not to be an open-ended exploration system. We provide freedom for user-driven exploration by allowing for limited personalization and "detours" from the set path of a data tour (section 5, BALANCE). Detours are possible by inserting related facts into an ongoing tour, as suggested by NetworkNarratives' recommender engine. For example, following the most connected node in the Network Overview tour, (4th panel in Figure 1) the detour inserts two slides about the 2nd and 3rd most connected nodes in that (Figure 1(D)).

4 EXAMPLE DATA TOURS

NetworkNarratives currently implements 10 tours, whose goals and main facts are detailed as follows. The list is not meant to be exhaustive but show the richness of possible tours. The full list of all tours, including their facts, as well as some illustrations can be found online: https://networknarratives.github.io/tours.

Network Overview describes an entire network. It starts with four introductory slides, covering the geographic extent (skipped for non-geographic networks), number of nodes and links, and the density. The second section focuses on links, showing the total and average link weight and the strongest

and weakest links in the network. The third section provides some details about node centralities and the overall community structure (number of clusters).

Subgraph Overview is similar to Network Overview, but focuses on a specific subgraph. It shows the subgraph's size and the percentage of the network's nodes. The tour also comprises important nodes such as the *Most connected node* in the subgraph, *Subgraph density*, and important links to the rest of the network.

Community Exploration explores and compares clusters in the network and shows their sizes, connections, and important nodes. For example, the *Most connected cluster* displays the cluster that has the most connections with the others. For community detection, we use the algorithm by Newman [50]. Advanced community detection algorithms can easily be included and used for comparison (e.g., k-means with different values for k) and shown on different slides (e.g., one slide for each value of k).

Centrality Exploration explores nodes based on different centrality measures (e.g., degree or betweenness). For example, we compute the *Average degree centrality*, the node with the *Highest betweenness centrality*, and the node with the *Highest Closeness centrality*. Possible extensions include comparisons of several centrality measures.

Subgraph Comparison compares two specified subsets of nodes and links (e.g., regions, subgraphs). Selecting this tour prompts the user to select two sets of nodes. The data tour first mentions the *Number of nodes* and the *Number of links* for each subgraph, then details important nodes such as the *Most connected node* in each subgraph, and finally reveals links between the two subgraphs (e.g., *Number of links*, *Total link weight*, and *Strongest link*).

Compare Two Nodes shows the links between the two nodes, compares their connectivities and total link weights, and finally shows the common neighbors. For example, general statistics such as *Connectivity ranking* and *Total link weight* of the two nodes are compared. Neighboring nodes that connect both of the selected nodes are shown in the last slide (*Common neighbors*) of the tour.

Ego-Network explores the network around a selected node and its neighbors. The data tour starts with the selected node and its position within the entire network. The tour then shows the node's direct neighborhood (nodes, links, strong connections), followed by their mutual connections, and finally its neighbors' neighbors.

Possible Paths explores a set of possible paths between two selected nodes. The data tour reports the path length, combined weights along each path, and the minimum link weight within each path. This data tour is motivated by *Arch*'s desire to explore historical travel costs between cities.

Follow a Path requires a selection of a set of connected nodes. The tour follows the path, explaining details about each node and its neighbors, and provides overall statistics of all the nodes in the path. This tour is motivated by *Arch*'s interest in nodes along geographic features, such as rivers, main roads, or political boundaries.

Temporal Exploration starts with an overview, showing the basic statistics about the network. It then demonstrates the connectivity evolution of the network based on temporal attributes. The data tour ends with the comparison of *Network density* over different time periods.

5 NETWORKNARRATIVES USER INTERFACE

To explore the potential of our guided data tours and allow for real-world use, we built a web-based prototype system called *NetworkNarratives*. NetworkNarratives can visualize and create tours of geographic and temporal networks that may include multiple link types and link weights. Networks are shown on a map if geographic information is given. Otherwise, they are rendered by using a force-directed layout.

The user interface (Figure 2) of NetworkNarratives consists of six panels: (a) data upload and selection; (b) a list of available data tours; (c) a detailed outline of the content of the currently selected data tour; (d) the network visualization with (e) captions; and (f) a navigation panel. Below, we explain how a user applies NetworkNarratives to explore a dataset of commuter movements in the Netherlands [25].

5.1 Importing data and defining terminology

After import, the analyst is prompted to specify domain-specific terminology for nodes, links, link-weight, and subgraphs. To illustrate, in the commuter network example, they might refer to nodes as cities, links as flows, and link weight as number of commuters. This feature was added in response to one of our analysts and aims to help relate explained facts easily to the domain [17].

5.2 Choosing and scoping a data tour

A user can then choose any data tour from the selection panel (Figure 2(B)). Hovering the cursor over a data tour's name displays a tooltip containing a short description. This approach simplifies browsing and selecting tours of interest (LEARN, SURPRISE, REPEAT). Alternatively, the analyst can select nodes and subgraphs in the visualization view by clicking on nodes or using a lasso interaction and subsequently choosing a data tour exploring the selected subgraph or node (e.g., *Ego-Network Exploration*, *Subgraph Overview*, and *Subgraph Comparison*).

5.3 Starting a data tour

Once a tour is selected from the tour panel in Figure 2(A), its structure becomes visible in the outline panel (Figure 2(C)) showing a tour's sections and facts in a tree-view. The visualization view shows a popup with the title and description of the tour and prompts the user to start the tour. The example in Figure 1 starts with the *Network Overview* tour.

Clicking the "*Click to start*" button in the popup loads the first *slide* in visualization view. Each slide refers to a single fact in the network. The first row in Figure 1 shows selected screenshots from the Network Overview tour. A slide includes *i)* a title representative of the fact shown on that respective slide (e.g., "Number of Nodes"), *ii)* an interactive visualization of the network (Figure 2(D)), and *iii)* a caption stating the corresponding fact (e.g., *"This network has*

11,216 links."). Any specific nodes or links mentioned by a fact are highlighted in the visualization.

Whenever a fact mentions a term that might be unfamiliar to the novice analyst or that might require additional explanation (e.g., link density), NetworkNarratives displays a hyperlink for a popup window. In the case of link density, the popup explains the formula used to calculate link density in NetworkNarratives (TRANSPARENCY, LEARN).

5.4 Navigating a data tour

An analyst can navigate to the next slide in the data tour by using the [▶ Next slide] button. A popup over the visualization indicates when the analyst has reached the second section of the tour. They can continue stepping through the slides by using the [▶ Next slide] button, or choose to skip that section and press the [▶▶ Next section] button. Alternatively, the analyst can return or jump directly to any specific slide in a section by clicking on its title in the outline view (Figure 2(C)).

In each slide, the network visualization is interactive to allow for exploration and the specification of nodes and links for potential tours (TRANSPARENCY). Within a visualization, an analyst can pan, zoom, and hover to highlight the connections of a specific node.

5.5 Detours

A detour inserts additional facts into a tour. To include a detour, an analyst can click the [⊕ More about [this slide]] button in the navigation panel. For example, if they are viewing the slide about the strongest link, then two new slides for the second and third strongest link are inserted (Figure 3(a)). Clicking the [⊕ More about [this slide]] button again appends additional relevant slides from NetworkNarratives' fact library that are recommended by the system (see section 6 for details). Likewise, detours can be included on a section level, i.e., to extend the current section. Clicking the [⊙ More about [the section]] button inserts slides with relevant facts about that section that are recommended by NetworkNarratives (Figure 3(b)).

Finally, the analyst can filter slides about facts that are not of interest. By clicking the "*filter*" button in the data tour selection panel, the analyst can uncheck any irrelevant aspects (e.g., *time* and *link weight*). Internally, these topics are stored in the form of *tags* associated with the facts and used to suggest related facts (section 6).

5.6 Pivoting to related data tours

Pivoting is possible for any slide discussing specific nodes or subgraphs. In these cases, NetworkNarratives displays a small popup that suggests to the user that they can pivot to a a related data tour (i.e., the yellow and blue tours in Figure 1). If the user decides to embark on a new tour, the new tour will be started (BALANCE). Likewise, NetworkNarratives suggests related tours at the end of each tour.

5.7 Personalizing and sharing data tours

The last option to introduce flexibility (BALANCE) is to personalize tours or create new tours. For example, slides of interest can be bookmarked in any tour using a star button. An analyst can edit
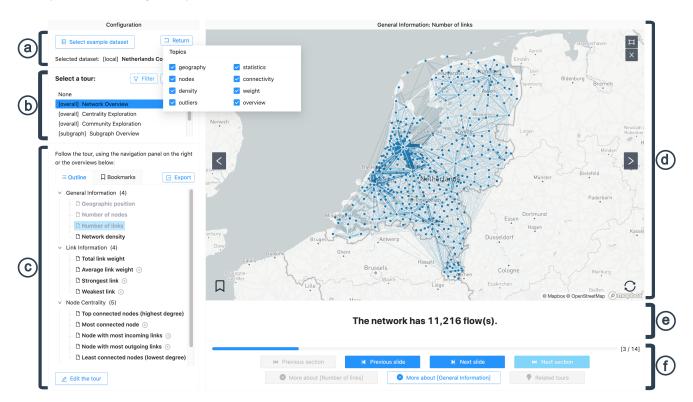
Figure 2: User interface of NetworkNarratives. The left sidebar contains the data selection panel (a), data tour selection panel (b), and tour overview and starred slides panel (c). Each slide consists of an interactive visualization (d) accompanied by a textual description (e). The user can move between slides by clicking on the arrow buttons in the visualization panel (d), using the buttons in the navigation panel (f), or clicking on the slide names in the outline panel (c).
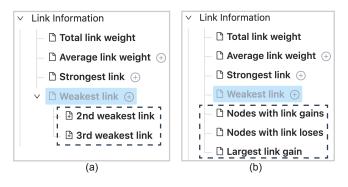

Figure 3: Slides and sections can be extended to show additional related slides. This example shows the results of extending (a) the *Weakest link* slide and (b) the *Link Information* section. In both images, the *Weakest link* slide is selected for display.

any data tour by clicking the "*Edit the tour*" button at the bottom of the data tour panel. In the editor panel, they can add new slides from NetworkNarratives' library of 102 facts. Selecting interesting tags (e.g., `weight`, `outliers`, and `link`) of the slides, filters the list of slides to display only those that are relevant for completing or generating new sections with a central topic. The analyst can also change the order of slides in a given tour, remove irrelevant

slides from a tour, or create entirely new tours from the fact library. Personalized tour templates can be exported in JSON format and shared with peers for reuse. Please refer to the supplementary video for additional details.

6 IMPLEMENTATION NOTES

NetworkNarratives is implemented as an open-source web-application. We used D3.js [15] for data processing and manipulation, flowmap.gl [70] and deck.gl [74] to render geospatial network visualizations, mapbox [43] to render the underlying map, The Vistorian library [4] to create interactive node-link diagrams, and React [47] to build the interactive interface.

Internal Tour Specification. In NetworkNarratives, tours are defined by a JSON specification, making the creation of new templates straightforward and flexible for developers. A specification defines a tour's id, name, and scope. The scope describes which network elements need to be defined by the analyst at the beginning of a tour: the overall network (nothing needs to be selected), a subgraph, or a single node. The specification then contains sections and their respective array of fact-IDs (slides). Examples of JSON specifications can be found on our website: https://networknarratives.github.io/tours.

Tagging Facts. NetworkNarratives has a library of 102 fact templates about networks, including information about the network,

subgraphs or individual nodes and links. Facts about centrality measures, topological information, attribute information, comparisons, rankings, outliers, paths, geographic information, clusters, or connectivity trends are included. A fact template is described by four attributes:

- **scope** describes what part of the network is targeted: the *overall* network, a set of *nodes*, or a single *node.*
- **link type** specifies whether a slide is suitable for networks with *directed* links, networks with *undirected* links, or *both.*
- **weight type** specifies whether the slide is suitable for networks with *weighted* links, networks with *unweighted* links, or *both.* Together, the `link type` and `weight type` define what networks a slide can be applied to.
- A fact can have one or multiple **tag**s (usually 2-3), which are used when selecting slides to recommend as detours. Tags are not meant to be exhaustive nor mutually exclusive. Tags for each fact were chosen by two of the paper's co-authors and currently include a total of 15 terms, such as `geography`, `nodes`, `links`, `weight`, `outliers`, `connectivity`, `statistics`, `density`, and `extrema`.

Fact Recommendation. When the user requests a detour, i.e., additional slides for a slide or section (section 5), NetworkNarratives identifies the most relevant slides to display. First, it counts the number of the occurrences of each topic tag in each section and calculates the TF-IDF score [59] as a measure of the importance of each tag to each section; this approach gives a higher weight to tags that are used more often in a section than to other tags, but decreases the weight of tags that are also used frequently in other sections. Each tour section is then represented by a vector, in which the i'th element is the TF-IDF score for that section and the i'th keyword. We calculate the cosine similarity between these vectors and return three randomly selected slides from the sections with the highest similarity to the section currently being expanded. The *link type* and *weight type* tags are used to exclude slides that could not be applied to the network being explored.

7 EVALUATING DATA TOURS

We performed two complementary user studies to understand the effectiveness of data tours in NetworkNarratives and the extent to which we achieved the design goals listed in subsection 3.2 (Learn, Reduce, Surprise, Repeat, Balance, Transparency). One study used experts in network analysis (subsection 7.1), whereas the other used novices (subsection 7.2). We report on the combined results in subsection 7.3.

7.1 Expert Evaluation

The expert study investigated the extent to which data tours could reduce an analyst's workload (Reduce, Repeat) and whether they could provide meaningful and potentially surprising insights (Surprise). Participants must have expertise in network analysis and an intrinsic interest in exploring their data to be able to compare data tours to their existing tools and workflows and to assess the usefulness of data tours. We invited all five analysts who were involved in the initial tour design (*Soc1, Soc2, Arch, Hist1, Health*) to individual exploration sessions with NetworkNarratives by using data provided by each analyst. We added three more analysts

who were not involved in the initial tour design: an analyst from academia (*Hist2*) and two senior developers from an industrial communication company with 5-7 years of experience in data development (*Dev1, Dev2*). During individual one-hour sessions, we first demonstrated NetworkNarratives and its main features and then provided each analyst with the tool URL, so they could run NetworkNarratives on their own machine and explore their own data. We recorded each session, asked the analysts to think aloud and ask for help when needed, and conducted interviews to solicit additional qualitative feedback about the tours and NetworkNarratives interface. Please refer to the supplementary material for the questions in the semi-structured interview.

7.2 Novice Evaluation

Our second study evaluated NetworkNarratives with network exploration novices. It focused on understanding whether tours help learn about network concepts (Learn) and the comprehensibility of the facts presented.

Conditions. We compared two conditions: **Tours** involved our NetworkNarratives interface with all tours, the navigation panel, and full interaction within each view (pan, zoom, mouse over, and selection). We removed the tour overview panel (Figure 3) to simplify the user interface for the study and because that view is not crucial to the concept of data tours. We compared Tours with the **Free-form** condition as a baseline because the novices would not know what to compare Tours with. Free-form showed only the network visualization from NetworkNarratives, including its interactions for exploration (pan, zoom, highlight, and selection). No tour or other navigation was available. Free-form was meant to be representative of any existing interactive free-form network visualization tool, such as Gephi [11] or The Vistorian [4].

Participants. We recruited 14 (seven females and seven males) postgraduate students (N1-N14) from a local university. All participants had a background in data science but did not have any experience in network exploration. No participant had prior knowledge of NetworkNarratives or other network exploration tools. We compensated the participants with a £10 gift card for their participation.

Setup and data. We ran a within-subject study with each participant experiencing both conditions. Half of the participants started with Tours, and the other half started with Free-form. For each condition, the participants received a 5-min introduction to the respective interactions, followed by 10-min exploration of the network or tours. We provided two geospatial network datasets—one for each condition—with weighted and directed links. One dataset was about internal migration in Sri Lanka. It comprised 25 cities (nodes), and 600 migration routes between cities (links). The other dataset was about bicycle hires in London. It contained 786 bike stations (nodes) and 10,000 trips (links) between these stations. We chose these two datasets because they are real-world datasets of similar complexity and do not require specific domain knowledge. All participants started with the Sri Lanka migration network independently of whether they started with Tours or Free-form. Each individual study lasted approximately 40 minutes.

After each condition, we asked the participants to explore the data to gain a comprehensive understanding and describe patterns
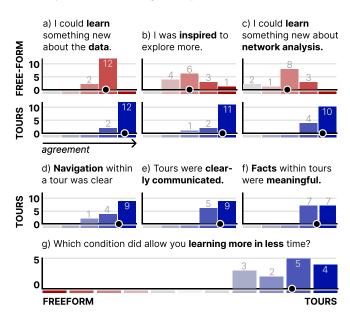
Figure 4: Subjective ratings from the novice study. Darker shades imply stronger agreement with the question posed. Dark circles indicate average values.

or insights that would be worth sharing with others. We asked the participants to take screenshots and note down 3-5 major findings in a 5-min time limit. We also encouraged participants to verbalize their thought process during exploration or when experiencing tours. We deliberately did not ask participants to complete a quiz because a fixed set of questions could have favored tours where facts are explicitly provided. On the other hand, tours would not have had a great chance in delivering knowledge that they did not cover. Throughout the session, the participants were allowed to seek assistance for the user interface whenever necessary and end the session early if they felt they had explored the data sufficiently.

Data collection. All sessions were held online. We recorded screens, think-aloud comments, interaction logs about the number of tours and facts visited, and the time spent on each fact. At the end of the study, the participants completed a questionnaire that asked for subjective ratings about each condition on 5-point Likert scales.

7.3 Study Results

This section reports the results of both studies: qualitative comments from the novices and experts, as well as subjective ratings from the novices. Generally, data tours (Tours) received higher subjective ratings (Figure 4) and more positive comments (Table 1) than the Free-form condition. Comments highlighted positive traits, such as *orientating* and *narrating*, *inspiring further exploration*, *leading to further discovery*, or *helping learning* about analysis and exploration. With Tours, participants browsed 31.1 slides (facts) on average, spending approximately 17 seconds per slide. Given a time limit of 10 minutes, the novices explored on average 2.29 tours on their own, spending 8:48 minutes with Tours, whereas they spent only 5:00 minutes with Free-form. When asked *"Which condition*

Cond.	Advantages	Disadvantages
Tours	Orients users with a narrative. Automatically provides facts. Can inspire deepened exploration. May lead to additional discoveries. Easy navigation. Saves time. Helps learn about analysis.	Can limit thinking and feel passive. Explanations need to be chosen carefully. Supports a rich set of views.
Free-form	Provides greater flexibility for exploration.	Harder to obtain deep insights or spot patterns with low prominence. Requires more time and effort to interact. Requires users to know where to look/have an exploration strategy.

Table 1: Reported advantages and disadvantages for Tours and Free-form across both studies.

did allow for learning more in less time," the novices responded with an average rating of +3.71 towards Tours on an 11-point scale (Free-form=-5, Tours=+5, Figure 4(g)). The original questionnaire showed numbers ranging from 1 (Free-form) to 11 (Tours) to avoid the potentially biasing -/+ notation.).

The novices commented that producing insights during free-form exploration (Free-form) was difficult: *"I only get the most obvious insights; [it was] hard to further exploration and get more insights"* [N10]. One novice missed guidance for a more in-depth exploration: *"I will be attracted by the most prominent patterns, but after examining these, I feel it hard for me to do further exploration"* [N14]. On the other hand, the novices found that Free-form provided more flexibility for exploration than tours and that tours could limit thinking, leading to *"passively accepting facts"* [N9]. While not entirely surprising, this observation is interesting because all free-form exploration features in Free-form (pan, zoom, highlighting, and selection) were also present in the Tours condition. We believe that this perception was due to a mindset that the study might have created because it was investigating two apparently opposing conditions. For the deployment of data tours, this situation could mean that the user interface needs to be very explicit about any free-form interaction capabilities. For example, tours should explicitly mention, explain, and encourage the use of free-form interaction with visualization.

In the following, we focus on data tours and NetworkNarratives in the novice and expert studies.

Tours and facts were perceived to be useful. The novices rated tours as well communicated (Figure 4(e), avg=4.71 on a 1-5 Likert scale) and found facts within the tours as meaningful (Figure 4(f), avg=4.5). The analysts concurred, commenting that these facts are *"in fact, what I would look at myself"* [Hist1], and that they are *"good for hypothesis generation"* [Hist1], *"cover[ing] a lot of common steps"* [Dev1] and *"guides what you should look at"* [Hist1]. For future improvements, the analysts suggested that the system could be extended by calculating additional network metrics, *"multiple attributes at the same time"* [Dev2], or *"templates that compare the layouts generated by different algorithms"* [Soc2, Dev1]. NetworkNarratives was designed to make such extensions straightforward. Likewise, the novices commented that the power of the approach grows with the number and content of the tours.

The participants highlighted the simplicity of exploring a network (*"I prefer NetworkNarratives because it directly guides me to see the next information"* [*Arch*]) (REDUCE) and that a rich set of views and information about the network is supported (*"provides multiple perspectives of the network."* [*Hist2*]), *"The slideshow analogy is helpful [...] I can see my time saving with this"* [*Arch*]). The novices suggested *"I like [data tours] because the data facts are organized systematically"* [N13] and *"I can read with little effort"* [N11]. Likewise, while not all facts in a given template might be of significant interest to an analyst, *Arch* expressed that facts which they would not have thought to request can nonetheless be very interesting, echoing another participant *"Helpful to know what I really want to look more into"* [*Soc1*] (BALANCE, SURPRISE).

NetworkNarratives can be integrated with existing workflows, e.g., *"to collect interesting facts and validate my observation[s]. And based on different basic facts, I still can tell different stories."* [*Hist2*]. While *"the existing data tours can satisfy most of my exploration needs"* [*Hist2, Dev1*], we can easily imagine additional tours to represent bespoke exploration methods and topics.

Tours have educational value and inspire exploration. The novices found that they could learn more about the data with TOURS (ratings in [Figure 4](a), TOURS=4.86, FREE-FORM=3.86), found TOURS to teach them more about network analysis ([Figure 4](c), TOURS=4.71, FREE-FORM=2.86), and were inspired to explore the data further ([Figure 4](b), TOURS=4.71, FREE-FORM=3.07). *"[Data tours] teach me how to analyze the network. Tours are like stories with different steps. I don't need to remember the key concepts. The network visualization explains well and clear[ly]"* [*Hist2*] (LEARN). Similarly, *Arch* suggested potential for educating students about network analysis, its concepts and methods and thought that the tours would be *"helpful to share with the new colleagues"* [*Arch*] for an introduction to network analysis. *Dev1* said *"I like the customizing and sharing functions because new employees can use the exported tour to get familiar with the data."* [*Dev1*] One participant highlighted the potential to generate new ideas and help develop a highly independent approach to network analysis: *"the information can be used to answer different questions and generate new ideas (e.g., findings or assumptions) or create a new story."* [*Hist2*] (SURPRISE). Another participant expressed a similar thought: *"the recommendation for querying more information around a specific topic would be useful when I have no idea about what story I could tell"* [*Dev2*] (SURPRISE). As highlighted by a recent study [4], this aspect is important for teaching goal and strategy development.

Navigation is easy. The user interface was perceived as simple and understandable, as reflected by the average rating of 4.64 given by the novices ([Figure 4](d)). The analysts and novices appreciated the simplicity of retrieving and displaying facts, especially network metrics, and that relevant nodes and links were highlighted in the visualization (*"The way of clicking the next button only is very friendly to me. I don't need other hints or reminder for what to do"* [*Hist2*], BALANCE). The participants could easily navigate forwards and backwards through the sequence of slides, making it *"easy to find back the information."* [*Hist2*]. At the same time, the analysts suggested a set of straightforward improvements. For example, *Hist2* suggested *"more instruction and tooltips"* [*Hist2*] and *Arch* recommended that each template could have a *"very quick"* preview, with the option to drill down and view additional details

on demand. We found this observation very interesting and adapted our design accordingly ([section 5](#)). Interestingly, N12 suggested providing further guidance on which tours to choose, especially when and why to pivot to a new tour.

Widened use cases. We discussed a range of potential applications for our approach in addition to personal exploratory analysis. *Hist1* said they would use the tool to make data and qualitative exploration available to peer researchers (*"It's frustrating because none of the publications provide interactions for the networks. Sharing data tours is valuable to play with by the others"* [*Hist1*]). *Dev1* and *Dev2* mentioned that *"the customized series of sequential fact is useful for reporting and presentation"* because they felt the interactive visualization was more vivid than the slides typically used for communication and demonstration within or outside their teams. *Arch* suggested potential for engaging with the public by using touch screens in museums, and is currently exploring this scenario for their own research. Finally, *Health* suggested using these tours to facilitate communication and discussions with policymakers.

8 DISCUSSION

Our study results provide strong evidence for the benefits of data tours and show that our design successfully supports our initial goals (G1-G6). The main findings can be summarized as follows:

(1) **Tours are an extensible concept.** Our current tours showed what our analysts were interested in. Yet, the sets of facts that can be shown in data tours are potentially very large, and the tour's power grows with the number of facts that they include. New facts can easily be added to our framework, which currently contains 102 individual facts.

(2) **Data tours are complementary to free-form tools.** The novices acknowledged that they found different insights under each condition (TOURS or FREE-FORM). They commented that data tours provide sufficient insights, but using them could be a passive activity. Thus, both conditions have their unique values to the users.

(3) **Data tours are a means to accelerate analysis and exploration** and reduce manual labor (REDUCE). Especially when dealing with numerous networks, the analysts are required to have consistency in exploration and analysis, as well as to explore these networks quickly (REPEAT).

(4) **Quick overview can prevent analysts from getting lost** in too many options and help them keep track of their previous exploration history, e.g., using a specialized tool [29]. For example, the analyst can overview a network through different tours, star interesting slides, and follow up on these slides in a second iteration.

(5) **Sequential tours support novice analysts** who are getting started with network visualization and learning about analysis methods and concepts (LEARN). They can be used by novices without previous knowledge about networks or specific network concepts or goals, who would otherwise struggle with a very open and free-form approach to network exploration [4]. Tours could help onboard and familiarize novice analysts with specific analysis routines, and even help engage and communicate networks and analysis to a broad audience.

(6) **Tours provide a serendipitous element** to exploration (SURPRISE). Comments from the user studies suggested that tours could also provide information quickly while keeping the door open for exploration.

Limitations. Evaluating exploration is an intrinsically difficult problem, given its open and context-sensitive nature. Our expert study aimed to account for this difficulty, while our novice study was limited in this regard. Although we attempted to choose datasets that might have been of interest to participants, finding datasets that are of equal interest to any participant in such a study while assuring that they are of equal complexity is difficult. We deliberately did not ask for specific insights (e.g., the number of clusters) to maintain the open nature of free-form exploration in FREE-FORM as a particular trait of that condition. We also acknowledge that data tours and free-form exploration are complementary and that future studies should compare data tours as human-designed and goal-oriented devices with pure recommender engines, such as DataShot [75]. However, at the time of our research, none of the existing recommender engines supported network analysis. Such a comparison could help assess to what extent these recommender systems are as transparent (TRANSPARENCY) as our data tours.

Although our study had only 10 tours available, we believe that they are representative of common network exploration tasks. We are now deploying NetworkNarratives as an add-on to The Vistorian [4] to understand data tours being used over a long period of time. We also aim to deploy NetworkNarratives to students of network analysis. However, such an effort would require further discussions with educators on how to best use NetworkNarratives in a curriculum.

Handling complex types of networks and analyses. Currently, NetworkNarratives can handle network data that are weighted or unweighted, geospatial or nongeospatial, and temporal. Although it does currently not support networks with multiple node and link types, adding respective facts to our slides would be trivial. Fully supporting temporal networks, however, appear to be a major challenge with respect to the tasks (facts) supported (e.g., [3]) and to the performance of meaningful temporal analysis to obtain facts [82]. The challenge is that complex types of networks also require complex analyses, for example, understanding and analyzing changes in dynamic networks [82]. Although task taxonomies [2] can be a starting point for design, we also need to better understand further strategies and workflows that analysts employ in practice to explore these complex types of networks.

Additional visual representations and presenting formats. Expanding the scope of NetworkNarratives to include complex types of networks and a richer collection of facts will require additional types of visualizations. This extension may include network visualizations [13, 61] or visualizations to summarize graph metrics, such as line graphs, bar charts, or ranked lists (e.g., [56]). Going beyond well-known visualization techniques and including a more diverse set of visual representations into tours will, in turn, require explanations of these techniques and their respective visual encodings (e.g., [78]). We are also investigating other ways of presenting networks, such as data comics [9] or videos [5] and the extent to which these formats foster engagement and understanding.

Overcoming challenges to automation. Automation can help extract a wide range of facts, such as network motifs, which we excluded from our current version for reasons of computational complexity. However, open questions as to the extent an algorithm can identify qualitative *insights* into networks remain. For example, we can calculate network metrics and topological features (motifs, paths, communities, and bridge nodes). However, how do we teach a machine to look for more subtle or complex patterns? Exploratory data analysis is powerful *because* it is done by humans.

Although machines can support this process by suggesting facts about data, automatically generating insights that incorporate users' domain knowledge and mental models will be difficult. One promising direction could be to begin by focusing on developing data tours for specific types of networks, such as biological pathways or social networks. These networks would have specific terminology and involve similar tasks, to which tour templates could be tailored.

Eventually, we imagine data tours being adjusted on-the-fly on the basis of the data being explored. For example, if a network has a "complex" community structure, a tour might provide additional information about that aspect. Likewise, if a network has many distinct communities, then a tour might report a few facts about each community. We believe that our linear, yet human-scripted, data tours are a first step toward complex tours that contain branches, loops, and other structural constructs. Our approach could inspire recommender systems that include further guidance and goal-oriented exploration. While understanding an analyst's analysis goals from interactions with a system is a well-known problem, data tours could propose directions then adapt to both data, and user interaction, while remaining direction.

Keeping data tours concise. Given that the collection of data facts and tour templates may continue to grow, keeping data tours concise will require active effort. Short and concise data tours were identified as an important criteria during our interviews. However, data tours need to contain sufficient facts to provide meaningful insights into networks.

We can address this situation in part by extending our data tours through appending additional related facts to sections or slides, as described in section 5. Additional approaches that could help keep data tours concise are available. One option is to split tours and create extended versions of the existing tours, which are designed *a-priori*. This approach might support expert analysts in performing complex network analyses. An alternative is to create direct sequels that pick up where a previous tour left off, but otherwise follow their own logic. Another option could be to use *starred* slides to create a user profile, akin to other existing systems (e.g., [65]). We believe that adding this technique to NetworkNarratives could be straightforward because it is conceptually independent from the dataset. Eventually, we could assign each fact in a tour a "priority" score and shorten or extend the tour on a user's demand (which could be expressed explicitly by using a simple slider, or inferred from user's interaction history by using a machine learning approach).

Support personalization, sharing, and storytelling. The growing number and size of possible data tours and facts suggests the potential for the personalization of tours and use of bespoke

storytelling approaches to maintain clarity and improve communication. The creation of personal templates in NetworkNarratives is currently supported through *starring* slides and editing tours. In the future, tour templates could be shared, commented on, modified by others, and reshared like software in the open-source community or plans for physical models in the maker community [31].

While NetworkNarratives is explicitly not designed for storytelling, data tours provide considerable potential for storytelling and presenting insights. Manually created data tours could be described as stories, especially if made by the network analyst themselves. However, it may require proper narrative language—in contrast to the currently rather factual explanations—that could include questioning, sign posting, analogies, metaphors, and other elements of narration to automate engaging storytelling [10]. Some of these extensions are a straightforward matter of scripting data tours in the form of personal stories and using story templates as suggested by *Arch*. We believe in the existence of an open space for future research to support storytelling about network data beyond current approaches, such as static data comics for networks [9, 35], or interactive comics that combine explanation and exploration [33, 77], slideshows [60], and data videos [6] for visualization. However, as illustrated by the large variety of narrative structures proposed in the literature, no universal recipe for creating good narratives exists. Storytelling is a very human activity that is inherently difficult to automate.

9 CONCLUSION

In this paper, we explored semi-automated data tours to aid network exploration and help learn concepts of network analysis. Our work is the first to explore data tours in the context of networks. Rather than implementing a full-fledged recommender engine, we opted to create goal-oriented tours inspired by real-world analysis practices. We see data tours as complementary to recommender approaches, as well as free-form exploration. At the same time, NetworkNarratives attempts to strike a balance between statically defined tours, data-driven recommender systems, and open-ended free-form exploration by including semi-automated techniques to retrieve related facts on user demand and allow for basic forms of personalization and exploration. We created an initial set of 10 data tours, and we imagine that additional data tours will be shared, modified, and reshared among analysts, potentially creating a global repository of tours. Feedback from expert analysts and novices suggests great potential for saving time and manual labor while helping to orient users. We believe that there is a promising future for future tools that can automatically identify and visually communicate insights by drawing on storytelling techniques.

ACKNOWLEDGMENTS

REFERENCES

[1] Guy J. Abel and Nikola Sander. 2014. Quantifying Global International Migration Flows. *Science* 343, 6178 (March 2014), 1520–1522. https://doi.org/10.1126/science.1248676

[2] Jae-wook Ahn, Catherine Plaisant, and Ben Shneiderman. 2014. A Task Taxonomy for Network Evolution Analysis. *IEEE Trans. Vis. Comput. Graph.* 20, 3 (2014), 365–376. https://doi.org/10.1109/TVCG.2013.238

[3] Jae-wook Ahn, Meirav Taieb-Maimon, Awalin Sopan, Catherine Plaisant, and Ben Shneiderman. 2011. Temporal Visualization of Social Network Dynamics: Prototypes for Nation of Neighbors. In *Social Computing, Behavioral-Cultural Modeling and Prediction - 4th International Conference, SBP 2011, College Park, MD, USA, March 29-31, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6589)*, John J. Salerno, Shanchieh Jay Yang, Dana S. Nau, and Sun-Ki Chai (Eds.). Springer, 309–316. https://doi.org/10.1007/978-3-642-19656-0_43

[4] Mashael AlKadi, Vanessa Serrano, James Scott-Brown, Catherine Plaisant, Jean-Daniel Fekete, Uta Hinrichs, and Benjamin Bach. 2023. Understanding Barriers to Network Exploration with Visualization: A Report from the Trenches. *IEEE Trans. Vis. Comput. Graph.* 29, 1 (2023), 907–917. https://doi.org/10.1109/TVCG.2022.3209487

[5] Fereshteh Amini, Nathalie Henry Riche, Bongshin Lee, Christophe Hurter, and Pourang Irani. 2015. Understanding Data Videos: Looking at Narrative Visualization through the Cinematography Lens. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18-23, 2015*. ACM, 1459–1468. https://doi.org/10.1145/2702123.2702431

[6] Fereshteh Amini, Nathalie Henry Riche, Bongshin Lee, Andrés Monroy-Hernández, and Pourang Irani. 2017. Authoring Data-Driven Videos with DataClips. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2017), 501–510. https://doi.org/10.1109/TVCG.2016.2598647

[7] Daniel Asimov. 1985. The Grand Tour: A Tool for Viewing Multidimensional Data. *SIAM J. Sci. Statist. Comput.* 6, 1 (Jan. 1985), 128–143. https://doi.org/10.1137/0906011

[8] Benjamin Bach, Nathalie Henry Riche, Roland Fernandez, Emmanoulis Giannisakis, Bongshin Lee, and Jean-Daniel Fekete. 2015. Networkcube: Bringing Dynamic Network Visualizations to Domain Scientists. https://www.microsoft.com/en-us/research/publication/networkcube-bringing-dynamic-network-visualizations-domain-scientists/ Posters Compendium of InfoVis 2015.

[9] Benjamin Bach, Natalie Kerracher, Kyle Wm. Hall, Sheelagh Carpendale, Jessie Kennedy, and Nathalie Henry Riche. 2016. Telling Stories about Dynamic Networks with Graph Comics. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016*. ACM, 3670–3682. https://doi.org/10.1145/2858036.2858387

[10] Benjamin Bach, Moritz Stefaner, Jeremy Boy, Steven Drucker, Lyn Bartram, Jo Wood, Paolo Ciuccarelli, Yuri Engelhardt, Ulrike Koppen, and Barbara Tversky. 2018. Narrative Design Patterns for Data-Driven Storytelling. In *Data Driven Storytelling*. Taylor Francis, Chapter 5, 107–133. https://www.microsoft.com/en-us/research/publication/narrative-design-patterns-for-data-driven-storytelling/

[11] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. 2009. Gephi: An Open Source Software for Exploring and Manipulating Networks. In *Proceedings of the Third International Conference on Weblogs and Social Media, ICWSM 2009, San Jose, California, USA, May 17-20, 2009*. The AAAI Press. http://aaai.org/ocs/index.php/ICWSM/09/paper/view/154

[12] Michael Baur. 2008. *visone - Software for the Analysis and Visualization of Social Networks*. Ph. D. Dissertation. Karlsruhe Institute of Technology. http://digbib.ubka.uni-karlsruhe.de/volltexte/1000010897

[13] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. 2017. A Taxonomy and Survey of Dynamic Graph Visualization. *Comput. Graph. Forum* 36, 1 (2017), 133–159. https://doi.org/10.1111/cgf.12791

[14] Michael Behrisch, Benjamin Bach, Michael Hund, Michael Delz, Laura von Rüden, Jean-Daniel Fekete, and Tobias Schreck. 2017. Magnostics: Image-Based Search of Interesting Matrix Views for Guided Network Exploration. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2017), 31–40. https://doi.org/10.1109/TVCG.2016.2598467

[15] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D³ Data-Driven Documents. *IEEE Trans. Vis. Comput. Graph.* 17, 12 (2011), 2301–2309. https://doi.org/10.1109/TVCG.2011.185

[16] Jeremy Boy, Louis Eveillard, Françoise Détienne, and Jean-Daniel Fekete. 2016. Suggested Interactivity: Seeking Perceived Affordances for Information Visualization. *IEEE Trans. Vis. Comput. Graph.* 22, 1 (2016), 639–648. https://doi.org/10.1109/TVCG.2015.2467201

[17] Jeremy Boy, Ronald A. Rensink, Enrico Bertini, and Jean-Daniel Fekete. 2014. A Principled Way of Assessing Visualization Literacy. *IEEE Trans. Vis. Comput. Graph.* 20, 12 (2014), 1963–1972. https://doi.org/10.1109/TVCG.2014.2346984

[18] Davide Ceneda, Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jörg Schulz, Marc Streit, and Christian Tominski. 2017. Characterizing Guidance in Visual Analytics. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2017), 111–120. https://doi.org/10.1109/TVCG.2016.2598468

[19] Davide Ceneda, Theresia Gschwandtner, and Silvia Miksch. 2019. A Review of Guidance Approaches in Visual Data Analysis: A Multifocal Perspective. *Comput. Graph. Forum* 38, 3 (2019), 861–879. https://doi.org/10.1111/cgf.13730

[20] Christopher Collins, Natalia V. Andrienko, Tobias Schreck, Jing Yang, Jaegul Choo, Ulrich Engelke, Amit Jena, and Tim Dwyer. 2018. Guidance in the human-machine analytics process. *Vis. Informatics* 2, 3 (2018), 166–180. https://doi.org/10.1016/j.visinf.2018.09.003

[21] Tarik Crnovrsanin, Isaac Liao, Yingcai Wu, and Kwan-Liu Ma. 2011. Visual Recommendations for Network Navigation. *Comput. Graph. Forum* 30, 3 (2011), 1081–1090. https://doi.org/10.1111/j.1467-8659.2011.01957.x

[22] Gabor Csardi and Tamas Nepusz. 2006. The igraph software package for complex network research. *InterJournal Complex Systems* 1695, 5 (2006), 1–9. http://igraph.sf.net

[23] Zhe Cui, Sriram Karthik Badam, Mehmet Adil Yalçin, and Niklas Elmqvist. 2019. DataSite: Proactive visual data exploration with computation of insight-based recommendations. *Inf. Vis.* 18, 2 (2019). https://doi.org/10.1177/1473871618806555

[24] Çağatay Demiralp, Peter J. Haas, Srinivasan Parthasarathy, and Tejaswini Pedapati. 2017. Foresight: Recommending Visual Insights. *Proc. VLDB Endow.* 10, 12, 1937–1940. https://doi.org/10.14778/3137765.3137813

[25] flowmaps.blue (based on data from Statistics Netherlands). 2020. Commuters in the Netherlands. https://docs.google.com/spreadsheets/d/1Oe3zM219uSfJ3sjdRT90SAK2kU3xIvzdcCW6cwTsAuc/edit#gid=592795568

[26] Stefan Gladisch, Heidrun Schumann, and Christian Tominski. 2013. Navigation Recommendations for Exploring Hierarchical Graphs. In *Advances in Visual Computing - 9th International Symposium, ISVC 2013, Rethymnon, Crete, Greece, July 29-31, 2013. Proceedings, Part II (Lecture Notes in Computer Science, Vol. 8034)*. Springer, 36–47. https://doi.org/10.1007/978-3-642-41939-3_4

[27] Camille Harris, Ryan A. Rossi, Sana Malik, Jane Hoffswell, Fan Du, Tak Yeon Lee, Eunyee Koh, and Handong Zhao. 2021. Insight-centric Visualization Recommendation. abs/2103.11297 (2021). arXiv:2103.11297 https://arxiv.org/abs/2103.11297

[28] Christopher G. Healey and Brent M. Dennis. 2012. Interest Driven Navigation in Visualization. *IEEE Trans. Vis. Comput. Graph.* 18, 10 (2012), 1744–1756. https://doi.org/10.1109/TVCG.2012.23

[29] Jeffrey Heer, Jock D. Mackinlay, Chris Stolte, and Maneesh Agrawala. 2008. Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation. *IEEE Trans. Vis. Comput. Graph.* 14, 6 (2008), 1189–1196. https://doi.org/10.1109/TVCG.2008.137

[30] Kevin Zeng Hu, Michiel A. Bakker, Stephen Li, Tim Kraska, and César A. Hidalgo. 2019. VizML: A Machine Learning Approach to Visualization Recommendation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*. ACM, 128. https://doi.org/10.1145/3290605.3300358

[31] Makerbot Industries. 2013. Thingiverse–Digital Designs for Physical Objects. https://www.thingiverse.com/

[32] Ilir Jusufi, Christian Klukas, Andreas Kerren, and Falk Schreiber. 2012. Guiding the interactive exploration of metabolic pathway interconnections. *Inf. Vis.* 11, 2 (2012), 136–150. https://doi.org/10.1177/1473871611405677

[33] Daye Kang, Tony Ho, Nicolai Marquardt, Bilge Mutlu, and Andrea Bianchi. 2021. ToonNote: Improving Communication in Computational Notebooks Using Interactive Data Comics. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*, Yoshifumi Kitamura, Aaron Quigley, Katherine Isbister, Takeo Igarashi, Pernille Bjørn, and Steven Mark Drucker (Eds.). ACM, 727:1–727:14. https://doi.org/10.1145/3411764.3445434

[34] Alicia Key, Bill Howe, Daniel Perry, and Cecilia R. Aragon. 2012. VizDeck: self-organizing dashboards for visual analytics. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2012, Scottsdale, AZ, USA, May 20-24, 2012*, K. Selçuk Candan, Yi Chen, Richard T. Snodgrass, Luis Gravano, and Ariel Fuxman (Eds.). ACM, 681–684. https://doi.org/10.1145/2213836.2213931

[35] Nam Wook Kim, Nathalie Henry Riche, Benjamin Bach, Guanpeng Xu, Matthew Brehmer, Ken Hinckley, Michel Pahud, Haijun Xia, Michael J. McGuffin, and Hanspeter Pfister. 2019. DataToon: Drawing Dynamic Network Comics With Pen + Touch Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*. ACM, 105. https://doi.org/10.1145/3290605.3300335

[36] Younghoon Kim, Kanit Wongsuphasawat, Jessica Hullman, and Jeffrey Heer. 2017. GraphScape: A Model for Automated Reasoning about Visualization Similarity and Sequencing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, May 06-11, 2017*. ACM, 2628–2638. https://doi.org/10.1145/3025453.3025866

[37] Stanford University Humanities + Design lab. 2022. *Palladio: Visualize complex historical data with ease.* https://hdlab.stanford.edu/palladio/about.

[38] Stanford University Humanities + Design lab. 2022. Testimonials. http://hdlab.stanford.edu/palladio/testimonials.

[39] Po-Ming Law, Alex Endert, and John T. Stasko. 2020. Characterizing Automated Data Insights. In *31st IEEE Visualization Conference, IEEE VIS 2020 - Short Papers, Virtual Event, USA, October 25-30, 2020*. IEEE, 171–175. https://doi.org/10.1109/VIS47514.2020.00041

[40] Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. 2006. Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI Workshop on BEyond time and errors: novel evaluation methods for information visualization, BELIV 2006, Venice, Italy, May 23, 2006*. ACM Press, 1–5. https://doi.org/10.1145/1168149.1168168

[41] Doris Jung Lin Lee, Himel Dev, Huizi Hu, Hazem Elmeleegy, and Aditya G. Parameswaran. 2019. Avoiding drill-down fallacies with *VisPilot*: assisted exploration of data subsets. In *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI 2019, Marina del Ray, CA, USA, March 17-20, 2019*. ACM, 186–196. https://doi.org/10.1145/3301275.3302307

[42] Jock D. Mackinlay. 1986. Automating the Design of Graphical Presentations of Relational Information. *ACM Trans. Graph.* 5, 2 (1986), 110–141. https://doi.org/10.1145/22949.22950

[43] Mapbox. 2022. mapbox/mapbox-gl-js: JavaScript library for vector maps on the Web. https://github.com/mapbox/mapbox-gl-js

[44] Roberto Martínez Maldonado, Vanessa Echeverría, Gloria Fernandez Nieto, and Simon Buckingham Shum. 2020. From Data to Insights: A Layered Storytelling Approach for Multimodal Learning Analytics. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, Regina Bernhaupt, Florian 'Floyd' Mueller, David Verweij, Josh Andres, Joanna McGrenere, Andy Cockburn, Ignacio Avellino, Alix Goguey, Pernille Bjøn, Shengdong Zhao, Briane Paul Samson, and Rafal Kocielnik (Eds.). ACM, 1–15. https://doi.org/10.1145/3313831.3376148

[45] Thorsten May, Martin Steiger, James Davey, and Jörn Kohlhammer. 2012. Using Signposts for Navigation in Large Graphs. *Comput. Graph. Forum* 31, 3 (2012), 985–994. https://doi.org/10.1111/j.1467-8659.2012.03091.x

[46] Hrim Mehta, Amira Chalbi, Fanny Chevalier, and Christopher Collins. 2017. Datatours: A data narratives framework. In *IEEE InfoVis 2017-IEEE Information Visualization conference*. 1–2.

[47] Inc. Meta Platforms. 2022. React: A JavaScript library for building user interfaces. https://reactjs.org

[48] Dominik Moritz, Chenglong Wang, Greg L. Nelson, Halden Lin, Adam M. Smith, Bill Howe, and Jeffrey Heer. 2019. Formalizing Visualization Design Knowledge as Constraints: Actionable and Extensible Models in Draco. *IEEE Trans. Vis. Comput. Graph.* 25, 1 (2019), 438–448. https://doi.org/10.1109/TVCG.2018.2865240

[49] Tamara Munzner. 2014. *Visualization Analysis and Design.* A K Peters. https://www.worldcat.org/oclc/994718755

[50] M. E. J. Newman. 2004. Fast algorithm for detecting community structure in networks. *Physical Review E* 69, 6 (Jun 2004). https://doi.org/10.1103/physreve.69.066133

[51] Dang Tuan Nhon and Leland Wilkinson. 2014. ScagExplorer: Exploring Scatterplots by Their Scagnostics. In *IEEE Pacific Visualization Symposium, PacificVis 2014, Yokohama, Japan, March 4-7, 2014*, Issei Fujishiro, Ulrik Brandes, Hans Hagen, and Shigeo Takahashi (Eds.). IEEE Computer Society, 73–80. https://doi.org/10.1109/PacificVis.2014.42

[52] Chris North. 2006. Toward Measuring Visualization Insight. *IEEE Computer Graphics and Applications* 26, 3 (2006), 6–9. https://doi.org/10.1109/MCG.2006.70

[53] Deok Gun Park, Mohamed Suhail, Minsheng Zheng, Cody Dunne, Eric D. Ragan, and Niklas Elmqvist. 2022. StoryFacets: A design study on storytelling with visualizations for collaborative data analysis. *Inf. Vis.* 21, 1 (2022), 3–16. https://doi.org/10.1177/14738716211032653

[54] Thomas Lin Pedersen. 2022. *tidygraph: A Tidy API for Graph Manipulation.* https://tidygraph.data-imaginist.com, https://github.com/thomasp85/tidygraph.

[55] Adam Perer and Ben Shneiderman. 2006. Balancing Systematic and Flexible Exploration of Social Networks. *IEEE Trans. Vis. Comput. Graph.* 12, 5 (2006), 693–700. https://doi.org/10.1109/TVCG.2006.122

[56] Adam Perer and Ben Shneiderman. 2008. Systematic yet flexible discovery: guiding domain experts through exploratory data analysis. In *Proceedings of the 13th International Conference on Intelligent User Interfaces, IUI 2008, Gran Canaria, Canary Islands, Spain, January 13-16, 2008*, Jeffrey M. Bradshaw, Henry Lieberman, and Steffen Staab (Eds.). ACM, 109–118. https://doi.org/10.1145/1378773.1378788

[57] Adam Perer and Ben Shneiderman. 2009. Integrating Statistics and Visualization for Exploratory Power: From Long-Term Case Studies to Design Guidelines. *IEEE Computer Graphics and Applications* 29, 3 (May 2009), 39–51. https://doi.org/10.1109/mcg.2009.44

[58] Catherine Plaisant, Jesse Grosjean, and Benjamin B. Bederson. 2002. SpaceTree: Supporting Exploration in Large Node Link Tree, Design Evolution and Empirical Evaluation. In *2002 IEEE Symposium on Information Visualization (InfoVis 2002), 27 October - 1 November 2002, Boston, MA, USA*. IEEE Computer Society, 57–64. https://doi.org/10.1109/INFVIS.2002.1173148

[59] Gerard Salton and Michael McGill. 1984. *Introduction to Modern Information Retrieval.* McGraw-Hill Book Company. https://www.worldcat.org/oclc/951258526

[60] Arvind Satyanarayan and Jeffrey Heer. 2014. Authoring Narrative Visualizations with Ellipsis. *Comput. Graph. Forum* 33, 3 (2014), 361–370. https://doi.org/10.1111/cgf.12392

[61] Sarah Schöttler, Yalong Yang, Hanspeter Pfister, and Benjamin Bach. 2021. Visualizing and Interacting with Geospatial Networks: A Survey and Design Space. *Comput. Graph. Forum* 40, 6 (2021), 5–33. https://doi.org/10.1111/cgf.14198

[62] Hans-Jörg Schulz, Marc Streit, Thorsten May, and Christian Tominski. 2013. Towards a Characterization of Guidance in Visualization. In *Poster at IEEE Conference on Information Visualization (InfoVis)*.

[63] Edward Segel and Jeffrey Heer. 2010. Narrative Visualization: Telling Stories with Data. *IEEE Trans. Vis. Comput. Graph.* 16, 6 (2010), 1139–1148. https://doi.org/10.1109/TVCG.2010.179

[64] Jinwook Seo and Ben Shneiderman. 2005. A rank-by-feature framework for interactive exploration of multidimensional data. *Inf. Vis.* 4, 2 (2005), 96–113. https://doi.org/10.1057/palgrave.ivs.9500091

[65] Danqing Shi, Xinyue Xu, Fuling Sun, Yang Shi, and Nan Cao. 2021. Calliope: Automatic Visual Data Story Generation from a Spreadsheet. *IEEE Trans. Vis. Comput. Graph.* 27, 2 (2021), 453–463. https://doi.org/10.1109/TVCG.2020.3030403

[66] Fabian Sperrle, Davide Ceneda, and Mennatallah El-Assady. 2023. Lotse: A Practical Framework for Guidance in Visual Analytics. *IEEE Trans. Vis. Comput. Graph.* 29, 1 (2023), 1124–1134. https://doi.org/10.1109/TVCG.2022.3209393

[67] Arjun Srinivasan, Steven Mark Drucker, Alex Endert, and John T. Stasko. 2019. Augmenting Visualizations with Interactive Data Facts to Facilitate Interpretation and Communication. *IEEE Trans. Vis. Comput. Graph.* 25, 1 (2019), 672–681. https://doi.org/10.1109/TVCG.2018.2865145

[68] Christina Stoiber, Davide Ceneda, Markus Wagner, Victor Schetinger, Theresia Gschwandtner, Marc Streit, Silvia Miksch, and Wolfgang Aigner. 2022. Perspectives of visualization onboarding and guidance in VA. *Vis. Informatics* 6, 1 (2022), 68–83. https://doi.org/10.1016/j.visinf.2022.02.005

[69] Christina Stoiber, Florian Grassinger, Margit Pohl, Holger Stitz, Marc Streit, and Wolfgang Aigner. 2019. Visualization onboarding: Learning how to read and use visualizations. (2019). https://doi.org/10.31219/osf.io/c38ab

[70] Teralytics. 2018. teralytics/flowmap.gl: Flow map drawing layer for deck.gl. https://github.com/teralytics/flowmap.gl

[71] Elaine G. Toms. 2002. Information interaction: Providing a framework for information architecture. *J. Assoc. Inf. Sci. Technol.* 53, 10 (2002), 855–862. https://doi.org/10.1002/asi.10094

[72] John W. Tukey. 1977. *Exploratory Data Analysis.* Addison-Wesley. https://www.worldcat.org/oclc/03058187

[73] Stef van den Elzen and Jarke J. van Wijk. 2013. Small Multiples, Large Singles: A New Approach for Visual Data Exploration. *Comput. Graph. Forum* 32, 3 (2013), 191–200. https://doi.org/10.1111/cgf.12106

[74] vis.gl. 2022. visgl/deck.gl: WebGL-powered visualization framework for large-scale datasets. https://github.com/visgl/deck.gl

[75] Yun Wang, Zhida Sun, Haidong Zhang, Weiwei Cui, Ke Xu, Xiaojuan Ma, and Dongmei Zhang. 2020. DataShot: Automatic Generation of Fact Sheets from Tabular Data. *IEEE Trans. Vis. Comput. Graph.* 26, 1 (2020), 895–905. https://doi.org/10.1109/TVCG.2019.2934398

[76] Yun Wang, Haidong Zhang, He Huang, Xi Chen, Qiufeng Yin, Zhitao Hou, Dongmei Zhang, Qiong Luo, and Huamin Qu. 2018. InfoNice: Easy Creation of Information Graphics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21-26, 2018.* ACM, 335. https://doi.org/10.1145/3173574.3173909

[77] Zezhong Wang, Hugo Romat, Fanny Chevalier, Nathalie Henry Riche, Dave Murray-Rust, and Benjamin Bach. 2022. Interactive Data Comics. *IEEE Trans. Vis. Comput. Graph.* 28, 1 (2022), 944–954. https://doi.org/10.1109/TVCG.2021.3114849

[78] Zezhong Wang, Lovisa Sundin, Dave Murray-Rust, and Benjamin Bach. 2020. Cheat Sheets for Data Visualization Techniques. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020.* ACM, 1–13. https://doi.org/10.1145/3313831.3376271

[79] Leland Wilkinson and Graham Wills. 2008. Scagnostics distributions. *Journal of Computational and Graphical Statistics* 17, 2 (2008), 473–491. https://doi.org/10.1198/106186008X320465

[80] Graham Wills and Leland Wilkinson. 2010. AutoVis: Automatic visualization. *Inf. Vis.* 9, 1 (2010), 47–69. https://doi.org/10.1057/ivs.2008.27

[81] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock D. Mackinlay, Bill Howe, and Jeffrey Heer. 2016. Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations. *IEEE Trans. Vis. Comput. Graph.* 22, 1 (2016), 649–658. https://doi.org/10.1109/TVCG.2015.2467191

[82] Liwenhan Xie, James O'Donnell, Benjamin Bach, and Jean-Daniel Fekete. 2020. Interactive Time-Series of Measures for Exploring Dynamic Networks. In *AVI '20: International Conference on Advanced Visual Interfaces, Island of Ischia, Italy, September 28 - October 2, 2020*, Genny Tortora, Giuliana Vitiello, and Marco Winckler (Eds.). ACM, 26:1–26:9. https://doi.org/10.1145/3399715.3399922

[83] Yalong Yang and Sarah Goodwin. 2019. What-Why Analysis of Expert Interviews: Analysing Geographically-Embedded Flow Data. In *IEEE Pacific Visualization Symposium, PacificVis 2019, Bangkok, Thailand, April 23-26, 2019.* IEEE, 122–126. https://doi.org/10.1109/PacificVis.2019.00022

[84] Vahan Yoghourdjian, Yalong Yang, Tim Dwyer, Lawrence Lee, Michael Wybrow, and Kim Marriott. 2021. Scalability of Network Visualisation from a Cognitive Load Perspective. *IEEE Trans. Vis. Comput. Graph.* 27, 2 (2021), 1677–1687. https://doi.org/10.1109/TVCG.2020.3030459

[85] Li Yu, Aidong Lu, William Ribarsky, and Wei Chen. 2010. Automatic Animation for Time-Varying Data Visualization. *Comput. Graph. Forum* 29, 7 (2010), 2271–2280. https://doi.org/10.1111/j.1467-8659.2010.01816.x

[86] Jian Zhao, Mingming Fan, and Mi Feng. 2022. ChartSeer: Interactive Steering Exploratory Visual Analysis With Machine Intelligence. *IEEE Trans. Vis. Comput. Graph.* 28, 3 (2022), 1500–1513. https://doi.org/10.1109/TVCG.2020.3018724