

# iTrace : Interactive Tracing of Cross-View Data Relationships

Abdul Rahman Shaikh  
ashaikh2@niu.edu  
Northern Illinois University  
DeKalb, Illinois, USA

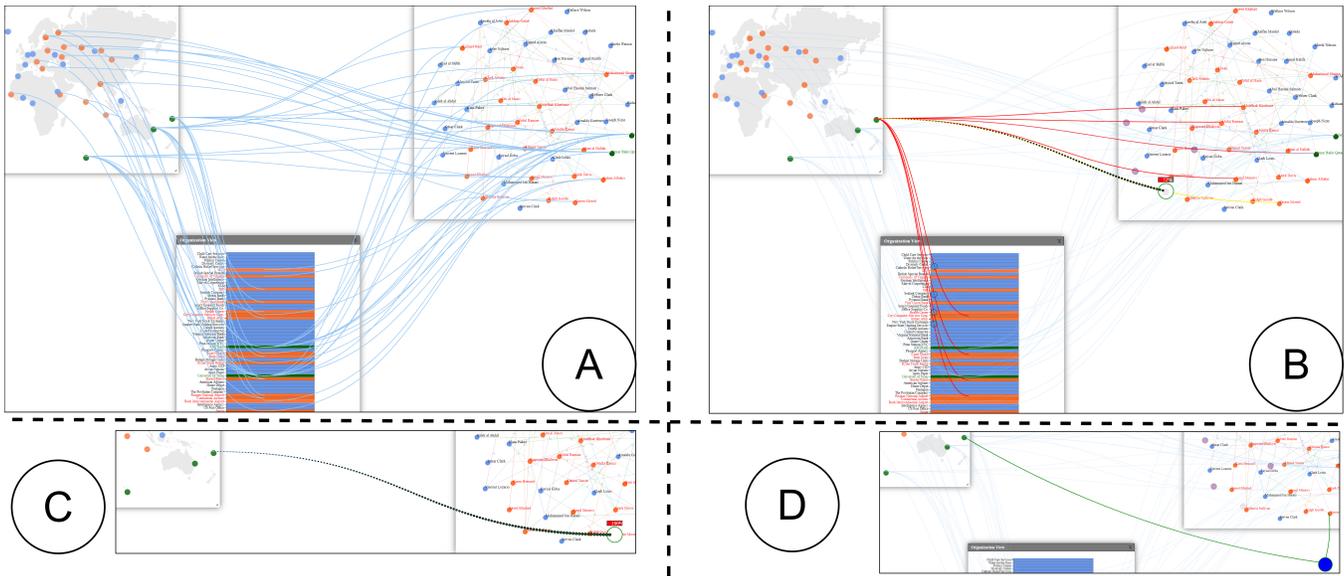
Hamed Alhoori  
alhoori@niu.edu  
Northern Illinois University  
DeKalb, Illinois, USA

Maoyuan Sun  
smaoyuan@niu.edu  
Northern Illinois University  
DeKalb, Illinois, USA

Jian Zhao  
jianzhao@uwaterloo.ca  
University of Waterloo  
Waterloo, Ontario, Canada

Xingchen Liu  
xliu7@niu.edu  
Northern Illinois University  
DeKalb, Illinois, USA

David Koop  
dakoop@niu.edu  
Northern Illinois University  
DeKalb, Illinois, USA



**Figure 1:** An example scenario using iTrace: (A) an initial stage before tracing, (B) tracing an element from a map to a graph, where supportive foci (circles on red lines) automatically follow their visual links, and dynamic transitions adjust other links to reduce clutter, (C) the final stage where the user finishes tracing, and (D) manual link management that to organize links.

## ABSTRACT

Exploring data relations across multiple views has been a common task in many domains such as bioinformatics, cybersecurity, and healthcare. To support this, various techniques (e.g., visual links and brushing & linking) are used to show related visual elements across views via lines and highlights. However, understanding the relations using these techniques, when many related elements are scattered, can be difficult due to spatial distance and complexity. To address this, we present iTrace, an interactive visualization technique to effectively trace cross-view data relationships. iTrace leverages the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Conference'17, July 2017, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

concept of *interactive focus transitions*, which allows users to see and directly manipulate their focus as they navigate between views. By directing the user's attention through smooth transitions between related elements, iTrace makes it easier to follow data relationships. We demonstrate the effectiveness of iTrace with a user study, and we conclude with a discussion of how iTrace can be broadly used to enhance data exploration in various types of visualizations.

## CCS CONCEPTS

• **Human-centered computing** → **Visualization techniques**.

## KEYWORDS

Multiple views, data relationship, visual analysis, interaction technique

### ACM Reference Format:

Abdul Rahman Shaikh, Maoyuan Sun, Xingchen Liu, Hamed Alhoori, Jian Zhao, and David Koop. 2025. iTrace : Interactive Tracing of Cross-View Data Relationships. In . ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

*Multiple-view visualization* (MV) is widely used for visual analysis in many fields (e.g., bioinformatics [23, 34, 39], cybersecurity [10, 80], finance [9, 31], health care [18], and education [11, 81]). To gain a deep understanding of the data, analysts often need to explore data relations across views [65, 67]. For example, cybersecurity analysts use network graphs to see device communications, lists to view suspicious organizations, and maps to check geographical locations, requiring them to explore relations across these views. To support such explorations, techniques like visual links and brushing & linking are heavily used. Visual links connect related elements across views, and brushing & linking highlight related elements after selections. However, the spatial distribution of related visual elements can make it hard to follow links or interpret highlights, especially when many connections overlap or when elements are scattered. Figure 2 shows an example of the challenge, which explores related visual elements in a graph from a selected location on a map.

Previous attempts to address this have used aggregation-based techniques to reduce the clutter of visual links [51, 60, 63–65, 67, 68]. Although decreasing the number of visual links reduces overlap, it does not eliminate the need for users to transition across views, as the layout of views remains unchanged. As users shift across views to explore cross-view relationships, they still rely on visual links. While fewer intersections make it easier to trace a single link, users must pay close attention to avoid losing track of the link. Any lapse in focus can disrupt their ability to realign with the original pathway. This challenge is amplified when tracing multiple links simultaneously. Without visual guidance, it is hard to navigate complex cross-view relationships effectively, even with fewer visual links. Techniques based on brushing & linking eliminate the need for visual links, thereby avoiding issues with line crossings. However, they struggle to explicitly outline the complex relationships between elements across views, especially after multiple selections. As each selection turns a group of elements in MVs into highlights, multiple brushings result in highlighting various groups. Relying only on highlights without visual links makes it difficult for users to discern relationships among elements across views.

To address these challenges, we introduce iTrace, an interactive visualization technique to facilitate tracing cross-view data relationships. Unlike conventional methods that leave users to piece together disjointed visual cues, iTrace features the design of *interactive focus transition*. This mechanism allows users to visibly transfer their “focus” from one view to another via an on-screen marker, allowing them to see, manipulate, and track their focus while navigating across views. Continuous visual guidance ensures that even when multiple links must be traced simultaneously, the active pathway remains clearly highlighted, preventing users from losing track of connections. Moreover, the versatility of iTrace makes it applicable to a wide range of visualizations, extending its usefulness beyond traditional multi-view scenarios. In summary, our contributions are threefold. First, we conduct an in-depth design analysis to support tracing cross-view data relationships, expanding the design space of MVs. Second, we introduce a novel design concept, the *interactive focus transition*, which emphasizes the visualization and direct manipulation of user focus points during

**Table 1: Four levels of cross-view data relationships [67].**

Relationship Level	Number of Views	Number of Visual Elements
<i>Individual level</i> (1 : 1)	2	2
<i>Group level</i> (1 : $m$ )	2	$1 + m$ ( $m > 1$ )
<i>Bi-group level</i> ( $m : n$ )	2	$m + n$ ( $m, n > 1$ )
<i>Multi-group level</i> ( $m : n : \dots : k$ )	$> 2$	$m + n + \dots + k$ ( $m, n, \dots, k > 1$ )

cross-view data exploration. Third, we develop a visualization prototype, iTrace, which embodies our design, and we evaluate its effectiveness through a user study.

## 2 RELATED WORK

### 2.1 Cross-View Data Relationship

Cross-view data relationships can be grouped into three categories (Figure 3). The first involves relationships *between visual elements* across views, focusing on the detailed connections between elements across views. It examines interactions within individual views and how elements from separate views are connected. For example, it explores how five people in a social network graph might be linked to six organizations in another view. This approach effectively “breaks” the boundaries of individual views to highlight components of potential knowledge graphs [32], linking disparate data elements scattered across views. The second category covers relationships *between visual elements and views*, offering a more abstract connection. Here, views add context or detail to visual elements in other views. For example, a line chart might show housing sales trends over the past six months for regions highlighted on a map. This relationship suggests a hierarchical structure among views, aligning with Shneiderman’s visual information-seeking mantra [57], where one view gives an overview and others provide details as needed. The third category looks at broader connections *between views*, prioritizing overall insights from collective data rather than individual elements. For example, documents placed together may indicate relevance, showing patterns that span multiple components. This approach highlights more conceptual relationships for deeper understanding [45].

This work focuses on those *between visual elements*, which require investigations of visual elements inside views. These low-level details across views form a key foundation for extracting higher-level insights, which are crucial to a data-driven sensemaking process [50]. The relationships can be categorized into four levels (Table 1), based on the cardinality of involved sets of visual elements from different views [67]. The complexity increases from the *individual* level to the *multi-group* level [65, 66].

### 2.2 Visualizing Data Relationships across Views

Three designs for visualizing cross-view data relationships are [65]: *connection*, *visual highlight*, and *spatial proximity*.

**Connection** uses visual links to show cross-view data relationships. It draws lines among related visual elements across different views, corresponding to a one-to-one cross-view data relationship. Examples include VisLink [12] and its variant versions [22, 60]. It is widely used in visual analysis tools (e.g., Bixplorer [21], MyBrush [38], and Flowstrates [6]). However, visual links can lead to clutter, making it hard for users to trace connections and understand more complex structures. To mitigate this, some studies have introduced aggregation techniques that group visual links into clusters (e.g.,

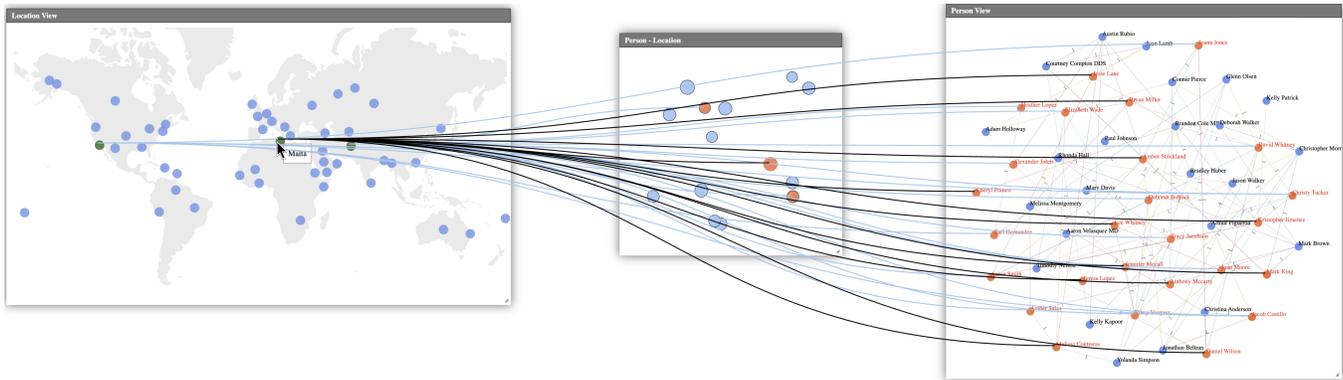


Figure 2: An example of the challenge on exploring related visual elements across views with visual links and highlights.

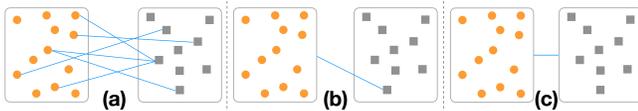


Figure 3: 3 types of cross-view data relationships: (a) *between visual elements*, (b) *between visual elements and views*, and (c) *between views*. Each box shows a view, orange circles and gray squares are visual elements, and blue lines indicate relations.

edge bundles) based on specific rules [51, 60, 64, 65, 67]. While these clusters simplify the presentation, users still need to trace individual lines to fully understand the underlying relationships. Notably, previous works have also addressed navigation and focus management across visual links. For instance, Moscovich’s link sliding and bring & go approaches [42] as well as Baudisch’s Drag-and-Pop/Drag-and-Pick technique [4] and the CompaRing method [69] similarly explore the concept of drawing copies of data entities to pull them toward the user’s focus.

**Visual highlight** relies on the coordination of MVs, with a dynamic update strategy to show relationships between visual elements across views. When users interact with elements in one view, corresponding elements in other views are highlighted. It is used in coordinated MVs [5, 46, 53], and is applied in various tools (e.g., Cross-Filtered Views [75], Improvise [74], Intervert [62], MissBiN [85], and Jigsaw [59]). While this technique helps users identify related elements, it requires continuous attention to visual updates triggered by user interactions. If these updates are missed, understanding the relationships can be hard. Unlike the connection approach, visual highlight avoids clutter by avoiding extra visual markers like lines, but without explicit links makes it hard to check complex cross-view relations.

**Spatial proximity** uses spatial distance to reveal relationships across views. It follows the “near equals similar” visual metaphor, which depends on the spatial arrangement of views to present their data relationships. This design requires users to interpret the spatialization of views to understand cross-view data relationships. It has been used in tools (e.g., Bixplorer [61], ForceSpire [20], GraphTrail [19], and NodeTriX [29]), which support “space to think” [2] oriented sensemaking activities. As spatial arrangement is controlled at the view level (i.e., manipulating the layout of MVs [56]) instead

Table 2: A Summary of Visualizations for Tracing

		Tracing Target (number of involved visual element)		
		One Element	Multiple Elements	
Analysis Focus (analytical space)	Data	One entity	One group	Multiple group
	Relation	- LineUp[24] - Parallel Tag Clouds [14] - Scatterplot Matrix [8]	- BubbleSets [13] - Crossets [47] - Kelpfusion[41] - LineSets [1] - Scatterplot Matrix [8]	- BiDot [84] - ChemoGraph [33] - Mercer [77] - MyBrush [38] - Sankey diagram [52] - VisLink[12]
	Location (spatial transition)	- SoccerStories [48] - Traffic flow visualizations [3] - Visual sedimentation [30]	- SoccerStories [48] - Traffic flow visualizations [3] - Trajectory bundling[17]	- Andromeda[55] - Animated transitions [28] - Traffic flow visualizations [3]
	Time (temporal trend, or process)	- A Table [49] - DimpVis [36] - egoSlider [79] - OpinionFlow [78] - StoryFlow [40] - ThemeRiver [27]	- EgoLines [82] - EventThread [25] - Matrixwave [83] - Story Curves [35] - VisTrails [7]	- EgoLines [82] - EventThread [25] - Reducing snapshots to points [72] - Matrixwave [83] - Story Curves [35]

of the visual element level, it is hard for users to identify and understand cross-view data relationships in detail. Moreover, interpreting spatializations requires more cognitive effort than tracking lines or visual highlights.

### 2.3 Visualizations for Tracing Information

Visualization techniques for tracing information focus on 4 aspects of visual analysis: 1) *data*, 2) *relation*, 3) *location*, and 4) *time* (Table 2). These mainly use *visual links*, *visual highlights*, and *animation*. Tracing involves individual or multiple visual elements that encode data from one or more groups. Such elements serve as tracing targets. Their encoded data corresponds to the analysis focus in the data space. This aligns with the information foraging loop in sensemaking [50]. The tracing-oriented analysis explores relationships, spatial transitions, and temporal trends based on the encoded data.

**Relation-focused tracing** involves analyzing connections between data entities (e.g., one-to-one, one-to-many, many-to-many). Simple cases involve tracing single visual elements [8, 14, 24] across different visual contexts (e.g., lists, scatterplots, small multiples [73]), using visual links or highlights. More complex cases trace multiple elements to identify those within the same group or across multiple groups. For the former, set visualization techniques [1, 13, 41] use lines or ribbons to connect elements of the same set, requiring users to follow them to trace visual elements of the same set. Brushing & linking based highlights are also used for tracing groups of

elements in different visual contexts (e.g., scatterplot matrix [8]). For the latter, visual links require users to follow multiple lines to trace elements across groups [33, 38, 77]. However, such techniques lack effective guidance, expecting users to manually trace paths.

**Location-focused tracing** examines spatial transitions by analyzing the movement of visual elements (e.g., sports or urban traffic analysis [3, 48]). Two main approaches are used: 1) line-based trajectories [3, 17, 48] where lines show paths, and 2) animations [17, 30, 55] showing dynamic transitions. Line-based trajectories allow users to follow movement patterns, and multiple lines can be bundled by spatial proximity. Animations guide user attention to moving elements, helping them stay focused. However, tracing animated elements is hard without visible trajectories, especially when elements move in different directions. Additionally, most techniques lack effective animation controls, forcing users to switch between traced elements and interface controls (e.g., play and pause), interrupting the tracing process and complicating task resumption.

**Time-focused tracing** involves temporal trends and process-oriented analysis, navigating through time to understand data evolution. Temporal trends are shown with lines [36, 49, 79] or areas [27, 78], where users follow shapes to observe value changes. For process-oriented analysis, stories, provenance, or event sequences are visualized using lines [40, 82], node-link diagrams [7, 72], and connected matrices [83]. Users follow line segments or paths to understand these processes. By adjusting spatial layouts (e.g., ordering lines [82] or matrices [83]), similar processes are placed near each other for comparison. However, users must manually trace paths without supportive features, which is hard with multiple processes.

In summary, while existing techniques use visual links, highlights, and animations, they lack assistance features and struggle to scale with multiple elements, especially for cross-view relationships. This motivates our design of iTrace.

## 3 DESIGNING ITRACE

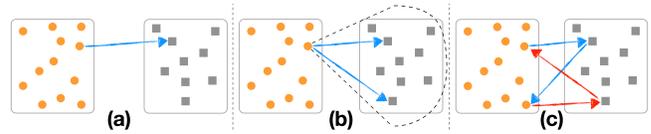
### 3.1 Term Clarification: Visual Link and Tracing

Our notion of **visual link** refers to a *perceptually continuous geometric shape* (e.g., a solid, dashed, or dotted line, curve, or ribbon) that connects related data scattered across different locations on a screen. These links are visual marks displayed over existing representations to explicitly connect related elements. This aligns with the definition of visual link discussed in [60], which highlights using additional visual marks instead of manipulating specific visual channels of existing visual marks (i.e., brushing and linking).

We consider **tracing** as an *analysis task* that guides users through conceptual spaces to find necessary information for sensemaking. For example, users can follow colored contours in a scatterplot to explore a *data space* and identify clusters, track edges in a graph to navigate a *relationship space* or observe line segments in a chart to examine a *time space* for trends. Specifically, in the context of cross-view data relations, tracing involves *following visual encodings to identify connections between visual elements across views*, which helps users understand relationships between visual elements.

### 3.2 User Tasks and Types of Tracing

In the context of cross-view data relationships, there are three major types of user tasks [67]: *filtering-oriented*, *refocusing-oriented*, and



**Figure 4: Three types of tracing: (a) individual oriented, tracing in a single direction; (b) group oriented, tracing in a constrained direction, and (c) cluster oriented, tracing in reflected directions. A blue/red arrow indicates tracing direction and dotted lines show a constrained range of tracing.**

*connecting-oriented* tasks. *Filtering-oriented* tasks refer to using visual elements in one view to filter those in other views. *Refocusing-oriented* tasks involve exploring the same data across different views based on selections made in one view. *Connecting-oriented* tasks aim to find connections between visual elements across views. These tasks correspond to data filtering, data identification, and data connection, respectively. As they all involve visual elements across views, tracing related elements serves as a key scaffold supporting the tasks. They are performed based on cross-view data relationships, particularly those between visual elements, as outlined in Section 2.1.

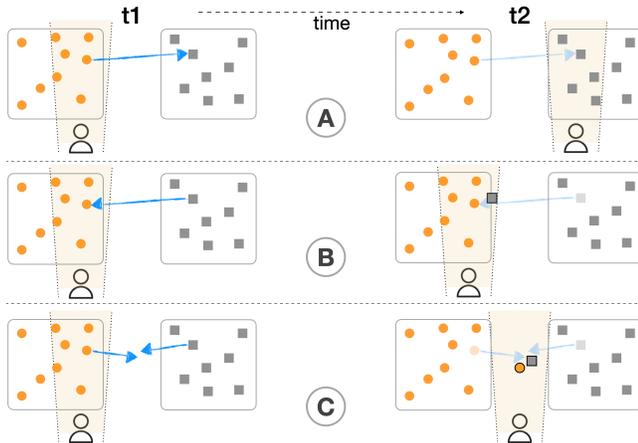
Considering the structure of cross-view data relationships (see Section 2.1), there are three primary types of tracing (see Figure 4): 1) *individual oriented* tracing, 2) *group oriented* tracing, and 3) *cluster oriented* tracing. These correspond to the first three levels of cross-view data relationships outlined in Table 1.

**Individual oriented tracing** ( $T1$ ) follows a single related visual element in a *single* direction, reflecting one-to-one relationships. **Group oriented tracing** ( $T2$ ) tracks relationships between a single element and multiple related elements within a bounded range, resulting in *constrained* directional tracing. Thus, it involves a group-level relationship (a one-to-many relationship). **Cluster oriented tracing** ( $T3$ ) examines relationships between two groups of elements, involving *reflected* directional tracing. This reveals a bi-group level relationship (a many-to-many relationship). While more complex relationships (multi-group level) exist, tracing them requires progressive analysis through multiple group or cluster-oriented traces, as they involve chains of information across multiple views. Therefore, this study focuses on the three primary types of tracing mentioned above.

### 3.3 Design Analysis: Trade-off & Consideration

Considering the change of a user's focus, as shown in Figure 5, there are three possible design strategies for supporting tracing: 1) *context switching*, 2) *context enriching*, and 3) *context separating*.

The **context switching** (S1) design strategy focuses on the necessity for users to transition their focus across views (Figure 5(A)). It acknowledges human cognition's limited capacity and the spatial separation of views, assuming users can only focus on specific elements within a single view at a time. While this approach preserves each view's original organization, it has notable drawbacks. The spatial separation of related elements and constant context switching increases cognitive load [15]. Users need clear visual cues to track related elements across views, especially when views are far apart. Without such cues, users can be confused about



**Figure 5: Designs for supporting tracing: context switching (A): user focus shifts from one view to another, context enriching (B): moving related visual elements from another view to the current one, and context separating (C): placing related visual elements outside original views and near each other. A trapezoid shows a user’s focus at a time. A blue arrow reveals a moving direction.**

which visual elements in the current view correspond to those previously checked.

The **context enriching** (S2) design aims to keep a user’s analysis focus within a single view by reducing the need to switch the focus between views (Figure 5(B)). This is achieved by moving related visual elements from other views into the current working view. This strategy enriches the visual context of the analysis, allowing users to trace cross-view data relationships without constant focus shifts, thereby reducing cognitive effort. However, this introduces new challenges. Moving elements between views requires reorganizing them, which can impact their structure. Users have to work on views in *transformative* forms, where visual elements in views are no longer confined within the traditional boundaries (i.e., visible borders) of a view. Instead, elements can be flexibly moved based on cross-view data relationships. This flexibility blurs boundaries, potentially complicating perception and leading to visual clutter.

The **context separating** (S3) design challenges the existing organizations of MVs by moving all related visual elements out of their original views (Figure 5(C)). It separates related visual elements from unrelated ones, effectively dividing a display space into two parts: 1) original views and 2) relationships among views. The former are the views as initially designed. The latter are areas where related elements are extracted into new types of views (e.g., relationship views). Such a division allows users to easily identify related visual elements, as they sit together and segregated from unrelated information, minimizing interference and visual clutter. Positioning related elements outside views can bring them into closer proximity than when dispersed across different views, minimizing the effort required for tracing due to reduced physical movement. However, this separation poses new challenges. Relocating visual elements strips them of their initial visual context, which might complicate users’ comprehension of cross-view data relationships.

Supporting tracing with these design approaches involves either redirecting a user’s focus between views or relocating visual

elements outside of their original views. The first strategy maintains the structure of views at the cost of requiring users to exert effort in transitioning between them. The second reduces user effort by minimizing the need to switch between views but disrupts the organization of views. The third allows users to easily identify connections without interference by segregating related elements from unrelated ones but strips elements of their original visual context. These compromises prompt a reevaluation of the standard organization of MVs.

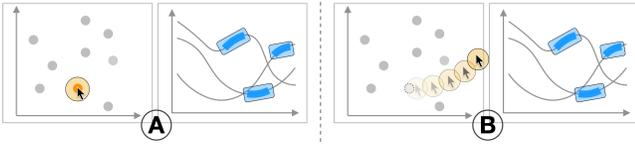
Based on our analysis of user tasks, types of tracing, and design strategies with trade-offs, we have identified four considerations that the design of iTrace attempts to follow.

**C1: Providing usable visual guidance to support transitioning between views.** Transitioning between views demands significant cognitive effort [15], which commonly occurs while using MVs, so it is vital to offer usable visual guidance to help users locate the target segment when moving from one view to another. Specifically, such visual guidance should direct a user’s attention to three key aspects of cross-view transitions involved in tracing: *direction*, *path*, and *destination*, corresponding to three main questions for tracing: which view(s) to navigate to, how to get there, and where to stop within the target view(s). (a) Showing the tracing *direction* helps users understand which view(s) they need to transition to. (b) Providing transition *paths* directs users’ attention during a transition process and helps to maintain focus. (c) Highlighting related visual elements within the target view(s) guides users to the *destination* elements they need to trace, effectively concluding the transition process.

**C2: Scaffolding multi-directional tracing.** As related visual elements are spatially scattered, it is important to offer techniques to support users in tracing multiple directions. Given the impact of selective attention [70], it is hard for users to trace several directions simultaneously, especially in divergent or opposite directions. While it is possible to turn this into a sequential process in which each step only involves tracing in one direction (**T1**), transitioning from a one-directional tracing to another creates challenges (e.g., which direction to pursue next and how users can ascertain completion of tracing in all directions). As multi-directional tracing (**T2**, **T3**) has been commonly involved in real-world analyses, developing techniques to facilitating this is imperative.

**C3: Comprehensive tracing of relationships.** To foster a transparent understanding of intricate interrelations among elements, it is crucial to allow users to navigate through all relationships associated with a visual element. In complex datasets, elements often have multiple connections spanning across elements in different views. Users should be able to confidently explore these connections. By supporting comprehensive tracing, users can efficiently trace individual elements (**T1**), follow connections between an element and a group (**T2**), and understand how groups relate to each other (**T3**). Without the ability to trace them, users may overlook critical connections, leading to incomplete analyses or incorrect conclusions.

**C4: Dynamic transition via tracing.** Dynamic transitions can be used to reduce visual complexity while a user is engaged in tracing an element. Factors influencing visual complexity include the number of elements, visual links, and views. Reducing visual



**Figure 6: Examples of externalizing a user's focus in iTrace. (A): iTrace creates a copy of each visual element involved in a cross-view data relationship and overlays them on existing views. (B): iTrace enables automatically moving a focus marker as a user moves a mouse pointer from its previously focused visual element. Orange circles and blue rectangles show the focus marker and copy of related visual elements in iTrace, respectively.**

complexity facilitates tracing and helps users effectively explore significant connections. Additionally, enhancing the salience of visual links supports users in tracing connections, allowing them to maintain focus and differentiate desired links. Informing users of their progress during tracing can improve their experience, aligning with usability heuristics (i.e., visibility of system status and recognition rather than recall) discussed by Nielsen and Molich [44]. Awareness of position and progress clarifies task completion, helping users see what has been explored and what remains. This feedback aids in resuming interrupted tasks and reduces cognitive load, especially when tracing spatially distant elements across multiple views.

## 4 THE ITRACE TECHNIQUE

The iTrace technique highlights a design concept of *interactive focus transition* to support cross-view data relationship tracing, which aims to make a user's focus visible and manipulable. It has three components: 1) externalizing a user's focus during tracing, 2) manipulating a user's focus to facilitate tracing, and 3) dynamic transition of a user's focus. They are applied on top of identified cross-view data relationships to support users in tracing them, so iTrace requires that cross-view data relationships are computed in advance. To compute the relationships, iTrace follows prior work [67], which uses *biclustering* to compute relationships between visual elements of pairwise views and chaining of biclusters (i.e., results of biclustering algorithms) to identify relationships among visual elements of three or more views. The computed cross-view data relationships serve as the foundation of iTrace, enabling it to provide visual cues and interactions that help users trace connections. They can be shown in various means, such as visual links connecting related elements, highlights that emphasize relationships, and dedicated relationship views that display visual elements together [67].

### 4.1 Externalizing User Focus

To make a user's focus explicit, iTrace introduces visual overlays, called **focus markers**, to show the user's current focus. Users can enable or disable a focus marker by right-clicking on a visual element. When enabled, a semi-transparent, slightly larger copy shows over the element (orange circles in Figure 6), distinguishing it from the original visual element without obscuring essential details. By interacting with multiple visual elements, users can create several focus markers simultaneously, allowing the tracing of bi-group level relationships (T3). Moreover, based on computed cross-view data relationships, iTrace enables users to generate copies of related

visual elements in other views (blue rectangles in Figure 6). These copies can also be enabled or disabled via a right-click menu on each focused visual element. They share visual cues with focus markers and represent *related elements*. This helps users recognize copies of visual elements, differentiate visual elements from their copies, and enable direct manipulations on focus markers.

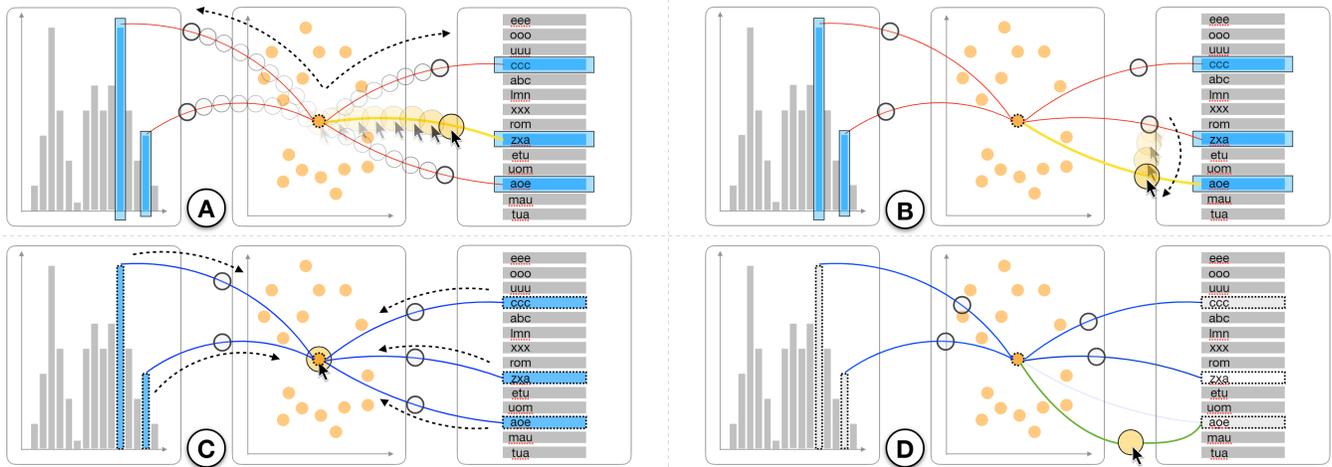
While the current implementation uses a semi-transparent, scaled copy for the focus marker, the iTrace technique is designed to be extensible, allowing its visual encoding to be customized for better visibility in complex scenarios, as noted in our user study. Instead of a direct copy, the marker could use distinct shapes (e.g., star, triangle, outlined circle), have static or density-aware sizing, or be enlarged relative to the original. Additional enhancements include high-contrast or saturated colors, stronger outlines or shadows, subtle animations (e.g., pulsing, blinking), or alternative fill styles like hatching. These adaptations can be tailored to view characteristics, information density, or user preference to ensure the marker remains easily distinguishable during tracing.

In iTrace, visual salience can propagate across multiple views based on computed cross-view data relationships, showing both a user's current and future focus. For example, in Figure 6, orange circles correspond to a user's current focus (i.e., where a tracing starts). Blue rectangles indicate the user's future focus (i.e., where the tracing will end). iTrace uses the visual salience in other view(s) to offer users guidance on tracing direction and destination (C1-a, c). After creating focus markers and copies, iTrace allows users to move these markers manually or automatically. For *automatic movement*, a focus marker moves by following a user's mouse pointer (Figure 6 (B)), externalizing the transition of the user's focus by transforming it from an internal cognitive process into visible on-screen elements. For *manual movement*, a user can drag and move the marker flexibly, which is discussed with more detail in Section 4.2.

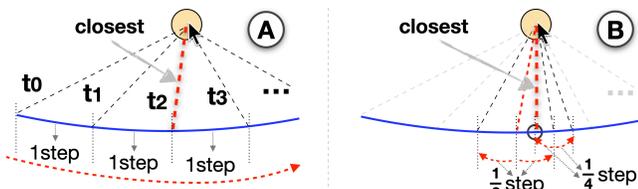
### 4.2 Manipulating User Focus

In iTrace, focus markers are considered the *first-level objects* for user interaction. Users can manipulate these markers by moving them to attract related visual elements from other views (Figure 7). Besides the automatic movement discussed in Section 4.1, users can manually move a focus marker by dragging it. To guide users in moving their focus marker, iTrace uses *visual links* (Figure 7 (A), (B)) from a visual element to its related elements, explicitly showing possible paths for users to perform tracing (C1-b). Users can move focus markers along these links (Figure 7 (A)). iTrace supports *cross-link transitions* with the focus marker (Figure 7 (B)). During the movement of a marker by a user, iTrace detects the closest point on any visual link connected to the current visual element and adjusts the marker's position accordingly. To find the closest point on a visual link, iTrace uses an algorithm with a linear [16] and iterative bidirectional search process [54]. Figure 8 shows the two procedures in this algorithm (see Algorithm 1). This algorithm is also used for automatic movement, ensuring that the focus marker "sticks" to the nearest visual link relative to the user's mouse pointer.

Along with the transition of a user's focus marker, iTrace creates **supportive foci** (white circles in Figure 7 (A) and (B)) that move along all other visual links associated with the element being traced.



**Figure 7: Examples of manipulating externalized user focus in iTrace. (A):** When a user moves the focus marker (i.e., a copy of a visual element in a scatterplot) along a visual link, supportive foci are displayed and automatically move along other visual links from the previously focused visual element. **(B):** A user moves the focus marker from one visual link to another. **(C):** From a currently focused visual element, a user calls for its related visual elements from other views, which shifts their copies along corresponding visual links. **(D):** A user interacts with a focus marker to manually manage links, positioning them as desired. Note: black dotted arrows reveal the direction of transitions, and the green link indicates a manually managed link.



**Figure 8: iTrace's two procedures to find the closest point on a visual link by the position of a user's focus: (A) a linear search process, and (B) an iterative bidirectional search process.**

These supportive foci can be enabled or disabled by left-clicking on the marker and are designed to assist in multi-directional tracing (C2) and facilitate comprehensive tracing of relationships (C3). This supports the context switching (S3) strategy, enabling users to shift their focus between related visual elements across multiple views. Supportive foci move in sync with the user's focus marker, with iTrace calculating their movement in real time based on the proportion of the marker's transition along its visual link. These supportive foci help users identify visual links that share the same cross-view data relationship as the focus marker. This reduces the effort required to locate relevant links, particularly in cases where many unrelated visual links are present. By acting as visual anchors, the supportive foci guide a user's attention to relevant connections. Research shows that humans better distinguish shapes (e.g., circles and lines) than color differences [43]. iTrace leverages this by using circles for supportive foci and lines for visual links. Moreover, when a user's focus marker transitions from one visual link to another, the supportive foci on the new link switch as well (Figure 7 (B)).

iTrace also allows the transition of related visual elements, using a *magnet and dust* visual metaphor [58]. From a focus marker (serving as the magnet, C1-a), the copy of its related visual elements (considered as the dust) can be attracted from other views (Figure 7 (C)). These copies move along the visual links that connect related

#### Algorithm 1: ClosestPoint(path, point)

```

Input : path, a path object; point, a point object
Output: closestPoint, the closest point object on a path
1 closestPoint, flagLen, closestDist  $\leftarrow$  none, 0,  $\infty$ ;
2 step, pathLen, scanLen  $\leftarrow$  8, path.getTotalLength(), 0;
3 while scanLen < pathLen do                                /* linear search */
4   scanPoint  $\leftarrow$  path.getPointAtLength(scanLen);
5   scanDist  $\leftarrow$  EuclideanDist(scanPoint, point);
6   if scanDist < closestDist then
7     | closestPoint, flagLen, closestDist  $\leftarrow$  scanPoint, scanLen, scanDist;
8   | scanLen  $\leftarrow$  scanLen + step;
9 while step > 0.5 do                                       /* bidirectional search */
10  step  $\leftarrow$  step / 2;
11  left, right  $\leftarrow$  flagLen - step, flagLen + step;
12  leftPoint, rightPoint  $\leftarrow$  path.getPointAtLength(left), path.getPointAtLength
    (right);
13  leftDist, rightDist  $\leftarrow$  EuclideanDist(leftPoint, point), EuclideanDist(rightPoint,
    point);
14  if leftDist < closestDist then
15    | closestPoint, flagLen, closestDist  $\leftarrow$  leftPoint, left, leftDist;
16  else if rightDist < closestDist then
17    | closestPoint, flagLen, closestDist  $\leftarrow$  rightPoint, right, rightDist;
18 return closestPoint;

```

elements to the focus marker (C1-b). Users can enable such transitions with a right-click menu on the element. iTrace offers two options in the right-click menu to control where the transitions stop: (C1-c): 1) at the border of the view where the focus marker is located (Figure 7 (C)), or 2) near the focus marker itself. The former aligns with the context-enriching design (S2) as user focus is maintained in a view. The latter supports the context-separating design (S3) when the focus marker is positioned away from the views. Users can hover over each moved copy and supportive foci to show information (e.g., an entity label) of its corresponding visual element in a popup. Moreover, when hovering over a copy, its corresponding visual mark in the original view is highlighted.

Moreover, iTrace introduces a *manual link management* feature that allows users to interact with visual links using focus markers

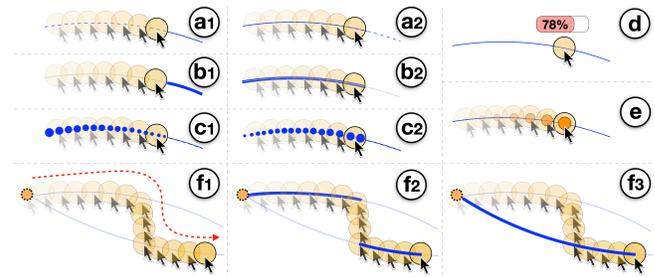
to drag and reposition them (Figure 7 (D)). Users can manage visual links manually, placing them anywhere within the display space. This is toggled via a right-click menu on the focus marker. Once activated, the focus marker moves along a user's mouse cursor dragging the visual link it was connected to at the time of activation. This feature allows users to store specific connections they may wish to explore later, effectively bookmarking relationships for future analysis. It also facilitates a more comprehensive tracing of relationships (C3) between two elements (T1) by providing users the flexibility to examine connections at their own pace. iTrace also enables the link management of multiple visual links. Users can make non-pinned links, either semi-transparent or hidden, via a right-click menu, allowing them to focus exclusively on the pinned connections. This enhances the flexibility of tracing, supporting both group-level tracing (T2) and bi-group level tracing (T3).

### 4.3 Dynamic Transition of User Focus

As a focus marker moves along a visual link, iTrace uses dynamic transitions to support comprehensive relationship tracing (C4). iTrace uses color encoding to indicate which link a user is focusing on. The visual link currently being traversed is highlighted in yellow, marking it as the active link. Meanwhile, related links connected to the same visual element, but not currently followed, are displayed in red, and all other unrelated links are colored blue. This color scheme allows users to distinguish among the active link, the related links available for further exploration, and the unrelated links.

iTrace uses *dynamic transparency control*, allowing users to adjust the visibility of unrelated links as they transition a focus marker. Users can choose to reduce the transparency of all unrelated links to focus solely on connections tied to the selected element or fade all links except the one being traced. The transparency adjusts dynamically with the transition progress. For instance, when a user reaches the midpoint of a link, unrelated links become 50% transparent, with opacity decreasing as the user advances. This progressive adjustment reduces visual complexity while preserving the link structure. Users can restore visibility by reversing transitions or revisiting previous paths. These provide users with more control over the display space, allowing for an in-depth exploration while reducing cognitive load. By highlighting relevant connections and subduing unrelated links, iTrace enhances user confidence in exploring relationships (C3). A challenge in designing dynamic transitions is maintaining effective visual emphasis in complex multi-view environments. We address this using a high-contrast color scheme and dynamic transparency, which highlights the active path and related options while fading unrelated links into the background. This reduces clutter and supports user focus, aligning with established foreground-background highlighting strategies in visualization [26].

As a focus marker moves, we explored visual aids to display the traced path, tracing direction, and transition progress (C1). Figure 9 shows our design alternatives. During the transition, the visual link is treated as two segments: the path traversed by the focus marker and the remaining part. Visual distinctions mark these segments: using dashed versus solid lines, varied line thickness, or dots of different sizes. While these designs indicate the traced path and direction, they do not effectively show the transition progress.



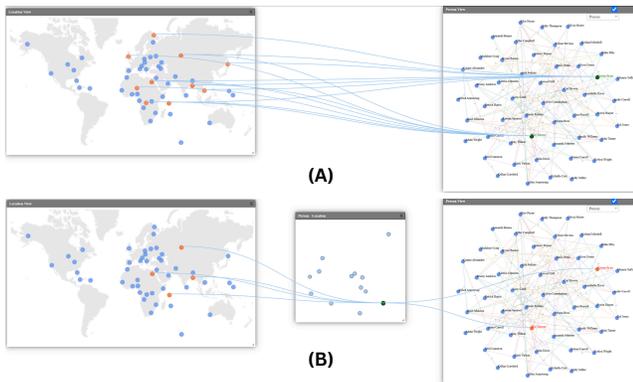
**Figure 9: Design alternatives to show a transition progress of a user's focus marker. (a2), (b2), and (c2) are encodings focusing on the path that a focus marker has passed. (a1), (b1), and (c1) show encodings emphasizing the remaining path. (d): Using a progress bar. (e): Using a proportional filling to show the progress (i.e., the more a focus marker has transitioned, the more filled area is inside the focus). For cases with shifts across links, three designs are considered: (f1) no path highlight, (f2) highlighting the path on each link that a focus marker has passed, and (f3) highlight on the current visual link only (to the most recent position of the focus marker), along which a focus marker is moving.**

Focus markers navigate links in two ways: 1) *single-link* transitions along one path (Figure 9 (a1, b1, c1)), and 2) *multiple-link* transitions between different paths (Figure 9 (f2)). For single-link transitions, the traced path, direction, and progress of a focus marker transition can be displayed by either adjusting the visual appearance of the visual link (Figure 9 (a1), (a2), (b1), (b2), (c1), and (c2)) or the focus marker itself (Figure 9 (e)), or adding additional marks (Figure 9 (d)). One design alternative uses *proportional filling* (Figure 9 (e)). It uses a circle within the focus marker that expands during transition. As a marker's size may limit the visibility of such changes, another alternative employs a *progress bar* design (Figure 9 (d)), where a bar is placed above the focus marker, offering better clarity in tracking movement compared to modifications of the link or marker alone.

For transitions with shifts across links, we explored three designs: 1) no path highlight, 2) highlighting all traversed links, and 3) highlighting only the active link up to a focus marker's position (Figure 9 (f1), (f2), and (f3), respectively). The first design shows only the marker position and gets ambiguous when multiple links intersect, making path tracking hard. The second design traces the complete transition path but creates visual clutter with numerous highlighted segments across multiple links. To balance clarity and usability, we adopted a third design that highlights only the current visual link from the trace origin to a focus marker's position. It reveals the transition process while maintaining clarity, even when other neighboring links are present. Our final design in iTrace highlights the current visual link being transitioned only (Figure 9 (f3)), accompanied by a progress bar positioned above the focus marker (Figure 9 (d)).

## 5 USER STUDY

To evaluate the usability of the key design concept, *interactive focus transition*, implemented in iTrace, we conducted a user study. Our study aimed to explore the effectiveness of this design concept for visual analysis involving tracing cross-view data relationships. We



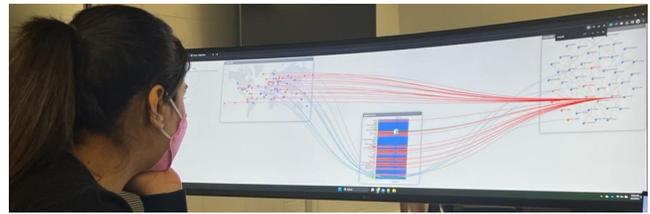
**Figure 10: Visualization of the bundling strategy. (A): Without bundling – individual visual links connect related entities across views, resulting in visual clutter. (B): With bundling – common relationships are aggregated and routed through a bundle element in a relationship view, simplifying the structure and revealing shared group-level connections.**

focused on three questions. First, would users prefer our interactive techniques over traditional MV techniques for tracing cross-view data relationships and why? This helps us understand the comparative appeal of our approach. Second, how do users apply our interactive techniques in practice? Observing their usage patterns offers insights into the intuitive use of our design. Finally, what do users perceive as the strengths and weaknesses of our techniques? This feedback is key to understanding the benefits and identifying areas for improvement.

### 5.1 Participants, Apparatus and Data

30 graduate students (14 female and 16 male) from several departments in a university, aged 23 to 33 years (mean age = 27), participated in our study voluntarily. All had normal or corrected-to-normal vision, no color vision deficiencies, and prior experience using multiple views. We deployed iTrace on a laptop with a 2.3 GHz Intel Core i7 processor and 16GB of memory, connected to a 49-inch ultrawide monitor with a resolution of  $5120 \times 1440$  pixels. Visualizations were shown in Google Chrome (version 116, 64-bit), which fit the screen without requiring scrolling. Participants sat approximately 55 cm from the monitor and used a mouse and keyboard for interaction. All participants provided informed consent prior to the study.

We generated eight datasets, each having three types of entities (e.g., person, location, and organization), with 50 entities of each type. Each type of entities was shown in a view: the person entities in a network graph, location entities on a map, and organization entities in a bar chart. We established cross-view data relationships by defining individual-level relationships between entities from different types (e.g., persons associated with locations). These relationships were set at two levels: by picking 10% or 20% of all pairwise entities (i.e., 2500 pairs) from two sets. Using these individual-level relationships, we applied LCM [71] to compute bi-group relationships between pairs of entity sets, ensuring that each group contained at least two entities. In summary, each dataset contains 3 different entity sets with a total of 150 unique entities, 250 or



**Figure 11: An example of the user study setting**

500 individual-level relationships, and 8~16 bi-group relationships. Figure 11 shows an example of the user study setting.

### 5.2 Method, Procedure and Data Collection

This study used a full-factorial, within-subjects design to explore the use of iTrace under varying conditions. We controlled two key factors: 1) **data complexity** and 2) **representation complexity**, as they influence the number of visible cross-view data relationships. For the former, we tested two levels: *fewer* versus *more* number of relationships in data (i.e., 250 vs. 500 individual-level relationships). For the latter, we set two levels: *without* versus *with* bundling visual links (i.e., grouping visual links that are in the same bi-group level relationships). As the two factors are independent, our study had four experiment conditions. In each condition, participants used three fixed-position visualizations: a map (left), a bar chart (center), and a network graph (right), each with 50 entities. In conditions with bundling, relationship views were displayed, grouping visual links that are in the same bi-group relationships as outlined in [67]. Specifically, if multiple elements in one view are commonly linked to elements in another view, their connections are aggregated and routed through an intermediate "bundle element" shown in a separate relationship view. This approach reduces clutter and emphasizes group-level associations by simplifying redundant links. The bundled connections are visualized explicitly to aid interpretation of shared relationships (Figure 10).

All of iTrace's techniques, described in Section 4, were available for participants to use in each condition. Each participant completed four information-foraging tasks (one task per experimental condition). The tasks were designed to reflect the three major types of tracing scenarios, corresponding to the levels of cross-view data relationships. An example task is: *"Find the organizations in North America where both the Andrews and Chandler families work."* This task requires participants to trace relationships between specific individuals (persons), locations, and organizations, involving both individual and group-level tracing. The order of tasks and their mapping to experimental conditions were randomized for each participant to counterbalance any learning effects.

Participants were given 15 minutes per task and could freely use any interactive techniques provided by iTrace based on their preferences. After each task, participants reported their findings and could take breaks as needed. Before the tasks, participants received a tutorial on cross-view data relations and multiple views, followed by a demonstration of iTrace's features using a separate dataset. A training session allowed them to practice identifying relationships across views. During each task, we collected interaction logs (e.g., timestamps, interaction types, and target elements), screenshots, observation notes, and participants' findings, and conducted interviews to gather their feedback. Accuracy was measured based on

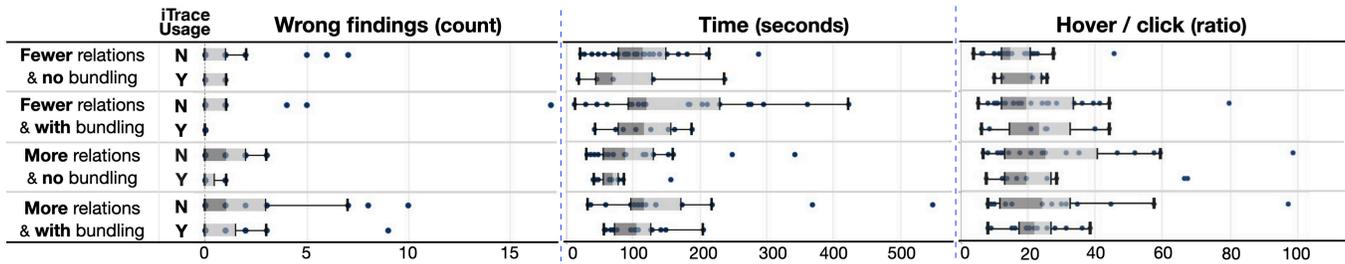


Figure 12: A summary of the number of participants' wrong findings (left), the average time that participants took to get a correct finding (middle), and the ratio between hover and click interactions that participants performed to find answers to the given tasks (right), in each experiment condition.

the correctness of participants' reported findings compared to the known relationships in the datasets.

### 5.3 Results

We observed that the number of participants using iTrace varied across conditions, and their task performance differed based on if iTrace was used (Figure 12). This suggests that the decision to use iTrace may be influenced by the perceived complexity of multiple views, affecting strategic interactions and leading to varied performance. Despite these variations, participants consistently got higher accuracy when using iTrace (Figure 13).

With the same visual link settings (i.e., with/without bundling), more participants chose to use iTrace when there were more cross-view data relationships. The increase in relationships led to more visual links. Participants may perceive this as visually overwhelming. For example, in conditions with more relationships, hovering over an entity would show a larger number of visual links than those with fewer ones. This increased visual complexity likely made it harder for participants to trace cross-view data relations, explaining the higher usage of iTrace in these scenarios. Moreover, more participants used iTrace when bundling was used than when it was not. Bundling reduced the number of visual links, but it introduced relationship views. While reducing visual links outweighed the increase in views, these new views added more visual elements (as each bundle was shown by an element), and visual links were re-routed through the bundles, making them longer. It made the representation more complex, encouraging more participants to rely on iTrace for assistance.

Regarding analysis performance, participants who used iTrace made fewer errors and took less time to find a correct answer than those without using iTrace (Figure 12 (left) and (middle)). This was consistent in all conditions. Even when the number of relationships or bundling settings changed, leading to variations in error rates and task completion time, participants using iTrace performed better overall. On average, participants using iTrace had an error rate of 0.5, whereas those not using iTrace had an error rate of 1.7. A Mann-Whitney U test indicated this difference was statistically significant ( $U = 180, p = 0.02$ ). Regarding task completion time, participants using iTrace took an average of 9.0 minutes, compared to 12.3 minutes for those not using it. An independent samples t-test confirmed this difference was significant ( $t(118) = 2.85, p = 0.005$ ).

The ratio of hover-to-click interactions was lower when iTrace was used than when it was not. Both hover and click interactions highlighted a visual element and its connected visual links. The

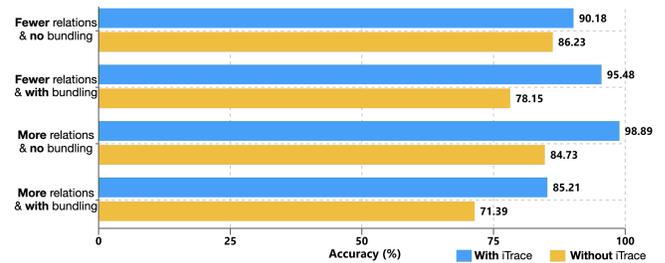


Figure 13: A summary of participants' finding accuracy.

highlights disappeared after hovering, while they remained after clicking. Thus, the hover-to-click ratio reflects how much effort participants put into observing changes in highlighting. Without iTrace, participants seemed to rely more on observing such highlight changes to trace cross-view data relations. When using iTrace, participants took less effort for this, as shown by the smaller hover-to-click ratio. This likely resulted from participants shifting their focus from observing highlight changes to actively manipulating visualized foci with iTrace. This redistribution of effort aligned with better analysis performance, as participants made fewer errors and took less time to find correct answers. Thus, iTrace appears to help users adopt more effective strategic interactions during analysis.

Besides these quantitative results and positive feedback on the "intuitive way of usage" (as noted by P6, P7, P9, and P13), participants' feedback highlighted three roles that iTrace played in supporting tracing cross-view data relations: 1) **scaling up tracing**, 2) **diversifying tracing purpose**, and 3) **strengthening confidence**. iTrace helped participants scale up their tracing capabilities by allowing them to perform multi-directional tracing even in the presence of many visual links. For example, P10 mentioned "It's really helpful to set up [see and use] copies to find a node connected to different other ones in other views"; P1 said, "[iTrace] is very useful especially when you have many things selected, and [it is] more useful when there are more connections"; and P17 appreciated, "It [iTrace] is very useful to find out the exact link when there is more data in the view." Additionally, iTrace expanded the purpose of tracing beyond just searching for connections. It also supported verification, helping participants confirm if two visual elements were connected. For example, P14 stated, "[iTrace] really helped to verify the data exactly and check appropriate connections... [iTrace] made the task easier to verify the connections"; and P7 said, "[iTrace] makes it easy to verify and understand the connections between data points." Finally, using

iTrace made participants more confident in their analyses. For instance, P6 said, “It helps me keep track of which line I’m looking at, so I’m more confident [of my findings]”, and P16 commented, “[iTrace] is good to give me the confidence to draw a conclusion.”

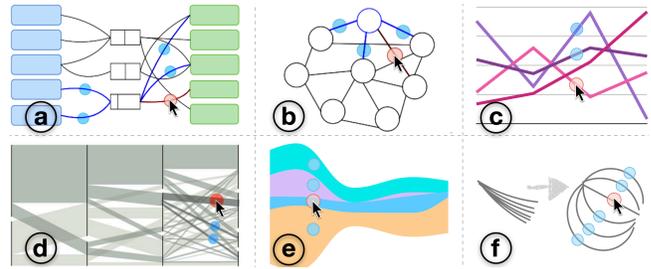
Despite the observed benefits of using iTrace, participants’ feedback indicated two possible improvements. First, while copies of visual elements created by iTrace can be helpful for tracing in multiple directions and find related elements quickly, it can sometimes lead to visual clutter. For example, P16 said, “It [iTrace] introduced more nodes [copies of visual elements] in the visualizations, and made it slightly confusing due to the visual clutter of many nodes on neighboring links”, and similarly P7 mentioned, “[copies of visual elements] get congested when there are too many connections.” These suggest that for visual elements with many visual links, it may be more effective to adjust the position of related visual element copies. Instead of placing them directly along the visual links, repositioning them to nearby areas could help reduce overlap and avoid visual clutter. Second, the focus marker’s visual encoding was perceived as insufficiently salient, especially when many visual links existed. For example, P15 mentioned, “When not moving [focus marker], it is hard to differentiate between its line and other lines.” P4 said similarly, “[After moving], it is confusing [locate where my focus was].” This implies that the current encoding of the focus marker lacks the visibility needed for participants to easily track their focus, particularly when there are many highlighted links and supportive foci.

## 6 DISCUSSION AND CONCLUSION

We present iTrace, an interactive technique to support tracing cross-view data relations. It introduces the concept of *interactive focus transition*, which transforms users’ focus into on-screen, manipulatable objects. By showing and allowing direct manipulation of these foci, iTrace facilitates multi-directional tracing across multiple levels of data relationships while reducing visual clutter associated with traditional linking methods. It enhances the flexibility and dynamism of MVs, moving beyond conventional, view-constrained layouts and instilling greater confidence in visual analysis.

### 6.1 Generalization

The *interactive focus transition* concept extends beyond MV oriented data analysis and can be adapted to a variety of visualizations (Figure 14). For instance, it can be used on node-link diagrams (e.g., graphs and trees) to check neighboring connections and paths, or to line charts for analyzing temporal trends across one or multiple lines. For visualizations with variant visual encodings of lines (e.g., ribbons in a Sankey diagram, parallel sets [37], or ThemeRiver [27]), this design concept can also be used. By assigning forces to supportive foci, the design can improve readability in cases with many overlapping lines by pulling them apart (Figure 14 (f)). This aligns with existing interactive techniques such as EdgeLens [76], which employs a lens-based interaction to effectively manage edge congestion in complex graphs. This directly addresses usability concerns related to visual clutter, as identified in Section 5.3. Compared to traditional visual links and highlights, this design concept makes the user’s focus more explicit by turning it into visible, on-screen objects that serve as anchors for attention and facilitate spatial transitions. Instead of replacing conventional techniques,



**Figure 14: Applying iTrace to other visualizations: (a) lists, (b) graph, (c) line chart, (d) parallel sets, (e) area graph, and (f) pulling nearby lines apart.**

it complements them by enhancing their ability to explore related information. Unlike typical animation techniques that rely on UI widgets (e.g., play and pause buttons), this design emphasizes direct manipulation. Users can interact with animated objects to control transitions, allowing them to focus on the content rather than the control mechanisms. This reduces cognitive load by eliminating the need to switch attention between a visualization and separate control widget(s).

### 6.2 Limitations and Future Work

The implementation of our presented *interactive focus transition* concept in iTrace supports direct manipulations on an individual copy only (e.g., focus marker or one of the supportive foci). It can be improved by grouping nearby foci moving in similar directions, enhancing iTrace for complex data handling. Future studies may explore how to merge and split these foci effectively. Some techniques, like manual link management and dynamic transitions in iTrace, introduce an extra level of control that requires further exploration. While our user study showed the usefulness of iTrace and its possible impact on tracing-based analysis, it was limited (e.g., datasets with a relatively small number of visual elements and views with fixed positions). Whether our findings would hold with larger datasets or varying distances between views remains uncertain, indicating the need for a more in-depth evaluation. Furthermore, participants’ self-selection to use iTrace during tasks led to unequal group sizes, limiting the ability to definitively attribute performance differences solely to the tool and highlighting the need for future studies with controlled usage conditions.

Additionally, iTrace potentially “breaks” traditional view boundaries by allowing flexible arrangements of visual elements and their copies to support tracing tasks. Future studies are needed to explore the potential costs of this approach, such as how users perceive and interact with these “broken” views. Performance issues such as lag and choppy interactions during tracing were caused by the real-time rendering demands of dynamic transitions and manipulations involving multiple visual links. These limitations can disrupt smooth tracing and negatively affect user experience. Although this study prioritized demonstrating conceptual and practical utility, future implementations could enhance usability by optimizing rendering through caching, incremental loading, or precomputed transitions.

### ACKNOWLEDGMENTS

This work was supported in part by the NSF Grant IIS-2002082.

## REFERENCES

- [1] Basak Alper, Nathalie Riche, Gonzalo Ramos, and Mary Czerwinski. 2011. Design study of linesets, a novel set visualization technique. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2259–2267. <https://doi.org/10.1109/TVCG.2011.186>
- [2] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to think: large high-resolution displays for sensemaking. In *Proc. of Conference on Human Factors in Computing Systems*. 55–64. <https://doi.org/10.1145/1753326.1753336>
- [3] Gennady Andrienko, Natalia Andrienko, Wei Chen, Ross Maciejewski, and Ye Zhao. 2017. Visual analytics of mobility and transportation: State of the art and further research directions. *IEEE Transactions on Intelligent Transportation Systems* 18, 8 (2017), 2232–2249. <https://doi.org/10.1109/ITITS.2017.2683539>
- [4] Patrick Baudisch, Ed Cutrell, Daniel Robbins, Mary Czerwinski, Peter Tandler, Benjamin Bederson, and Alex Zierlinger. 2003. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-Operated Systems. In *Human-Computer Interaction-INTERACT '03*. 57–64.
- [5] Nadia Boukhelifa, Jonathan Roberts, and Peter Rodgers. 2003. A coordination model for exploratory multiview visualization. In *Proceedings of International Conference on Coordinated and Multiple Views in Exploratory Visualization*. IEEE, 76–85. <https://doi.org/10.1109/CMV.2003.1215005>
- [6] Ilya Boyandin, Enrico Bertini, Peter Bak, and Denis Lalanne. 2011. Flowstrates: An approach for visual exploration of temporal origin-destination data. In *Computer Graphics Forum*, Vol. 30. <https://doi.org/10.1111/j.1467-8659.2011.01946.x>
- [7] Steven Callahan, Juliana Freire, Emanuele Santos, Carlos Scheidegger, Claudio Silva, and Huy Vo. 2006. VisTrails: visualization meets data management. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 745–747. <https://doi.org/10.1145/1142473.1142574>
- [8] Daniel B Carr, Richard J Littlefield, WL Nicholson, and JS Littlefield. 1987. Scatterplot matrix techniques for large N. *Journal of American Statistical Association* 82, 398 (1987), 424–436. <https://doi.org/10.1080/01621459.1987.10478445>
- [9] Remco Chang, Mohammad Ghoniem, Robert Kosara, William Ribarsky, Jing Yang, Evan Suma, Caroline Ziemkiewicz, Daniel Kern, and Agus Sudjianto. 2007. Wirevis: Visualization of categorical, time-varying data from financial transactions. In *IEEE Sym. on Visual Analytics Science and Technology*. IEEE, 155–162. <https://doi.org/10.1109/VAST.2007.4389009>
- [10] Siming Chen, Shuai Chen, Natalia Andrienko, Gennady Andrienko, Phong H Nguyen, Cagatay Turkay, Olivier Thonnard, and Xiaoru Yuan. 2018. User behavior map: Visual exploration for cyber security session data. In *IEEE Symposium on Visualization for Cyber Security*. <https://doi.org/10.1109/VIZSEC.2018.8709223>
- [11] Xin Chen, Jessica Self, Leanna House, John Wenskovich, Maoyuan Sun, Nathan Wycoff, Jane Evia, and Chris North. 2017. Be the data: Embodied visual analytics. *IEEE Tran. on Learning Technologies* 11, 1 (2017), 81–95. <https://doi.org/10.1109/TLT.2017.2757481>
- [12] Christopher Collins and Sheelagh Carpendale. 2007. VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1192–1199. <https://doi.org/10.1109/TVCG.2007.70521>
- [13] Christopher Collins, Gerald Penn, and Sheelagh Carpendale. 2009. Bubble sets: Revealing set relations with isocontours over existing visualizations. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1009–1016. <https://doi.org/10.1109/TVCG.2009.122>
- [14] Christopher Collins, Fernanda B Viegas, and Martin Wattenberg. 2009. Parallel tag clouds to explore and analyze faceted text corpora. In *IEEE Sym. on Visual Analytics Science and Technology*. 91–98. <https://doi.org/10.1109/VAST.2009.5333443>
- [15] Gregorio Convertino, Jian Chen, Beth Yost, Y-S Ryu, and Chris North. 2003. Exploring context switching and cognition in dual-view coordinated visualizations. In *Proceedings of International Conference on Coordinated and Multiple Views in Exploratory Visualization*. IEEE, 55–62. <https://doi.org/10.1109/CMV.2003.1215003>
- [16] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to algorithms*. MIT press.
- [17] Fan Du, Nan Cao, Jian Zhao, and Yu-Ru Lin. 2015. Trajectory bundling for animated transitions. In *Proceedings of the Conference on Human Factors in Computing Systems*. 289–298. <https://doi.org/10.1145/2702123.2702476>
- [18] Fan Du, Catherine Plaisant, Neil Spring, and Ben Shneiderman. 2016. EventAction: Visual analytics for temporal event sequence recommendation. In *IEEE Conference on Visual Analytics Science and Technology*. IEEE, 61–70. <https://doi.org/10.1109/VAST.2016.7883512>
- [19] Cody Dunne, Nathalie Henry Riche, Bongshin Lee, Ronald Metoyer, and George Robertson. 2012. GraphTrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. In *Proc. of Conf. on Human Factors in Computing Systems*. 1663–1672. <https://doi.org/10.1145/2207676.2208293>
- [20] Alex Endert, Patrick Fiaux, and Chris North. 2012. Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 473–482. <https://doi.org/10.1145/2207676.2207741>
- [21] Patrick Fiaux, Maoyuan Sun, Lauren Bradel, Chris North, Naren Ramakrishnan, and Alex Endert. 2013. Bixplorer: Visual analytics with biclusters. *Computer* 46, 8 (2013), 90–94. <https://doi.org/10.1109/MC.2013.269>
- [22] Thomas Geymayer, Markus Steinberger, Alexander Lex, Marc Streit, and Dieter Schmalstieg. 2014. Show me the invisible: visualizing hidden content. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3705–3714. <https://doi.org/10.1145/2556288.2557032>
- [23] Samuel Gratzl, Nils Gehlenborg, Alexander Lex, Hanspeter Pfister, and Marc Streit. 2014. Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2023–2032. <https://doi.org/10.1109/TVCG.2014.2346260>
- [24] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. 2013. LineUp: Visual analysis of multi-attribute rankings. *IEEE Trans. on Visualization & Comp. Graphics* 19, 12 (2013), 2277–2286. <https://doi.org/10.1109/TVCG.2013.173>
- [25] Shunan Guo, Ke Xu, Rongwen Zhao, David Gotz, Hongyuan Zha, and Nan Cao. 2017. EventThread: Visual summarization and stage analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 56–65. <https://doi.org/10.1109/TVCG.2017.2745320>
- [26] K Wm Hall, Charles Perin, Peter G Kusalik, Carl Gutwin, and Sheelagh Carpendale. 2016. Formalizing emphasis in information visualization. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 717–737.
- [27] Susan Havre, Elizabeth Hetzler, Paul Whitney, and Lucy Nowell. 2002. Themeriver: Visualizing thematic changes in large document collections. *IEEE Trans. on Vis. and Computer Graphics* 8, 1 (2002), 9–20. <https://doi.org/10.1109/2945.981848>
- [28] Jeffrey Heer and George Robertson. 2007. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1240–1247. <https://doi.org/10.1109/TVCG.2007.70539>
- [29] Nathalie Henry, Jean-Daniel Fekete, and Michael J McGuffin. 2007. NodeTriX: a hybrid visualization of social networks. *IEEE Transactions on Vis. and Computer Graphics* 13, 6 (2007), 1302–1309. <https://doi.org/10.1109/TVCG.2007.70582>
- [30] Samuel Huron, Romain Vuillemot, and Jean-Daniel Fekete. 2013. Visual sedimentation. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2446–2455. <https://doi.org/10.1109/TVCG.2013.227>
- [31] Dong Hyun Jeong, Wenwen Dou, Heather Richter Lipford, Felesia Stukes, Remco Chang, and William Ribarsky. 2008. Evaluating the relationship between user interaction and financial visual analysis. In *IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 83–90. <https://doi.org/10.1109/VAST.2008.4677360>
- [32] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and S Yu Philip. 2021. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems* 33, 2 (2021), 494–514. <https://doi.org/10.1109/TNNLS.2021.3070843>
- [33] Bharat Kale, Austin Clyde, Maoyuan Sun, Arvind Ramanathan, Rick Stevens, and Michael Papka. 2023. ChemoGraph: interactive visual exploration of the chemical space. In *Com. Graphics Forum*, Vol. 42. 13–24. <https://doi.org/10.1111/cgf.14807>
- [34] Bharat Kale, Maoyuan Sun, and Michael E Papka. 2023. The state of the art in visualizing dynamic multivariate networks. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 471–490.
- [35] Nam Kim, Benjamin Bach, Hyejin Im, Sasha Schriber, Markus Gross, and Hanspeter Pfister. 2017. Visualizing nonlinear narratives with story curves. *IEEE Trans. on Vis. and Computer Graphics* 24, 1 (2017), 595–604. <https://doi.org/10.1109/TVCG.2017.2744118>
- [36] Brittany Kondo and Christopher Collins. 2014. DimpVis: Exploring time-varying information visualizations by direct manipulation. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2003–2012. <https://doi.org/10.1109/TVCG.2014.2346250>
- [37] Robert Kosara, Fabian Bendix, and Helwig Hauser. 2006. Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics* 12, 4 (2006), 558–568. <https://doi.org/10.1109/TVCG.2006.76>
- [38] Philipp Koytek, Charles Perin, Jo Vermeulen, Elisabeth André, and Sheelagh Carpendale. 2017. Mybrush: Brushing and linking with personal agency. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 605–615. <https://doi.org/10.1109/TVCG.2017.2743859>
- [39] Alexander Lex, Marc Streit, Ernst Kruijff, and Dieter Schmalstieg. 2010. Caleydo: Design and evaluation of a visual analysis framework for gene expression data in its biological context. In *IEEE Pacific Visualization Symposium*. IEEE, 57–64. <https://doi.org/10.1109/PACIFICVIS.2010.5429609>
- [40] Shixia Liu, Yingcai Wu, Enxun Wei, Mengchen Liu, and Yang Liu. 2013. Storyflow: Tracking the evolution of stories. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2436–2445. <https://doi.org/10.1109/TVCG.2013.196>
- [41] Wouter Meulemans, Nathalie Riche, Bettina Speckmann, Basak Alper, and Tim Dwyer. 2013. Kelfusion: A hybrid set visualization technique. *IEEE Trans. on Vis. and Com. Graphics* 19, 11 (2013). <https://doi.org/10.1109/TVCG.2013.76>
- [42] Tomer Moscovich, Fanny Chevalier, Nathalie Henry, Emmanuel Pietriga, and Jean-Daniel Fekete. 2009. Topology-aware navigation in large networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. 2319–2328. <https://doi.org/10.1145/1518701.1519056>
- [43] Tamara Munzner. 2014. *Visualization analysis and design*. CRC press.

- [44] Jakob Nielsen and Rolf Molich. 1990. Heuristic evaluation of user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 249–256. <https://doi.org/10.1145/97243.97281>
- [45] 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>
- [46] Chris North and Ben Shneiderman. 1997. *A taxonomy of multiple window coordination*. Technical Report.
- [47] Charles Perin and Pierre Dragicevic. 2014. Manipulating multiple sliders by crossing. In *Proceedings of the 26th Conference on l'Interaction Homme-Machine*. 48–54. <https://doi.org/10.1145/2670444.2670449>
- [48] Charles Perin, Romain Vuillemot, and Jean-Daniel Fekete. 2013. SoccerStories: A kick-off for visual soccer analysis. *IEEE Transactions on Vis. and Computer Graphics* 19, 12 (2013), 2506–2515. <https://doi.org/10.1109/TVCG.2013.192>
- [49] Charles Perin, Romain Vuillemot, and Jean-Daniel Fekete. 2014. A Table! improving temporal navigation in soccer ranking tables. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 887–896. <https://doi.org/10.1145/2556288.2557379>
- [50] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of International Conference on Intelligence Analysis*, Vol. 5. 1–6.
- [51] Arnaud Prouzeau, Antoine Lhuillier, Barret Ens, Daniel Weiskopf, and Tim Dwyer. 2019. Visual link routing in immersive visualisations. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*. 241–253. <https://doi.org/10.1145/3343055.3359709>
- [52] Patrick Riehmann, Manfred Hanfler, and Bernd Froehlich. 2005. Interactive sankey diagrams. In *IEEE Symposium on Information Visualization*. IEEE, 233–240. <https://doi.org/10.1109/INFVIS.2005.1532152>
- [53] Jonathan C Roberts. 2007. State of the art: Coordinated & multiple views in exploratory visualization. In *Proceedings of International Conference on Coordinated and Multiple Views in Exploratory Visualization*. IEEE, 61–71. <https://doi.org/10.1109/CMV.2007.20>
- [54] Stuart J Russell and Peter Norvig. 2020. *Artificial intelligence: A modern approach, 4th edition*. Pearson. 82 pages.
- [55] Jessica Zeitz Self, Michelle Dowling, John Wenskovich, Ian Crandell, Ming Wang, Leanna House, Scotland Leman, and Chris North. 2018. Observation-level and parametric interaction for high-dimensional data analysis. *ACM Transactions on Interactive Intelligent Systems* 8, 2 (2018), 1–36. <https://doi.org/10.1145/3158230>
- [56] Abdul Rahman Shaikh, David Koop, Hamed Alhoori, and Maoyuan Sun. 2022. Toward systematic design considerations of organizing multiple views. In *IEEE Vis. and Visual Analytics*. 105–109. <https://doi.org/10.1109/VIS54862.2022.00030>
- [57] Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of IEEE Symposium on Visual Languages*. IEEE, 336–343. <https://doi.org/10.1109/VL.1996.545307>
- [58] Ji Soo Yi, Rachel Melton, John Stasko, and Julie A Jacko. 2005. Dust & magnet: multivariate information visualization using a magnet metaphor. *Information Visualization* 4, 4 (2005), 239–256. <https://doi.org/10.1057/palgrave.ivs.9500099>
- [59] John Stasko, Carsten Gorg, Zhicheng Liu, and Kanupriya Singhal. 2007. Jigsaw: supporting investigative analysis through interactive visualization. In *IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 131–138. <https://doi.org/10.1109/NAIST.2007.4389006>
- [60] Markus Steinberger, Manuela Waldner, Marc Streit, Alexander Lex, and Dieter Schmalstieg. 2011. Context-preserving visual links. *IEEE Trans. on Visualization & Computer Graphics* 17, 12 (2011), 2249–2258. <https://doi.org/10.1109/TVCG.2011.183>
- [61] Maoyuan Sun, Lauren Bradel, Chris L North, and Naren Ramakrishnan. 2014. The role of interactive biclusters in sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1559–1562. <https://doi.org/10.1145/2556288.2557337>
- [62] Maoyuan Sun and Gregorio Convertino. 2016. Interver: Drilling into categorical-numerical relationships. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*. 44–51. <https://doi.org/10.1145/2909132.2909253>
- [63] Maoyuan Sun, David Koop, Jian Zhao, Chris North, and Naren Ramakrishnan. 2019. Interactive bicluster aggregation in bipartite graphs. In *2019 IEEE Visualization Conference*. IEEE, 246–250. <https://doi.org/10.1109/VISUAL.2019.8933546>
- [64] Maoyuan Sun, Peng Mi, Chris North, and Naren Ramakrishnan. 2015. Biset: Semantic edge bundling with biclusters for sensemaking. *IEEE Trans. on Vis. & Computer Graphics* 22, 1 (2015), 310–319. <https://doi.org/10.1109/TVCG.2015.2467813>
- [65] Maoyuan Sun, Akhil Namburi, David Koop, Jian Zhao, Tianyi Li, and Haeyong Chung. 2021. Towards systematic design considerations for visualizing cross-view data relationships. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2021), 4741–4756. <https://doi.org/10.1109/TVCG.2021.3102966>
- [66] Maoyuan Sun, Chris North, and Naren Ramakrishnan. 2014. A five-level design framework for bicluster visualizations. *IEEE Transactions on Vis. and Computer Graphics* 20, 12 (2014), 1713–1722. <https://doi.org/10.1109/TVCG.2014.2346665>
- [67] Maoyuan Sun, Abdul Rahman Shaikh, Hamed Alhoori, and Jian Zhao. 2021. Sightbi: Exploring cross-view data relationships with biclusters. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 54–64. <https://doi.org/10.1109/TVCG.2021.3114801>
- [68] Maoyuan Sun, Jian Zhao, Hao Wu, Kurt Luther, Chris North, and Naren Ramakrishnan. 2018. The effect of edge bundling and seriation on sensemaking of biclusters in bipartite graphs. *IEEE Transactions on Visualization and Computer Graphics* 25, 10 (2018), 2983–2998. <https://doi.org/10.1109/TVCG.2018.2861397>
- [69] Christian Tominski. 2016. CompaRing: Reducing Costs of Visual Comparison. In *EuroVis 2016*, Enrico Bertini, Niklas Elmqvist, and Thomas Wischgoll (Eds.). The Eurographics Association. <https://doi.org/10.2312/eurovisshort.20161175>
- [70] Anne M Treisman. 1969. Strategies and models of selective attention. *Psychological review* 76, 3 (1969), 282–299. <https://doi.org/10.1037/h0027242>
- [71] Takeaki Uno, Tatsuya Asai, Yuzo Uchida, and Hiroki Arimura. 2004. An efficient algorithm for enumerating closed patterns in transaction databases. In *Discovery Science*. Springer, 16–31. [https://doi.org/10.1007/978-3-540-30214-8\\_2](https://doi.org/10.1007/978-3-540-30214-8_2)
- [72] Stef van den Elzen, Danny Blaas, Jorik Blaas, and Jarke J van Wijk. 2015. Reducing snapshots to points: A visual analytics approach to dynamic network exploration. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 1–10. <https://doi.org/10.1109/TVCG.2015.2468078>
- [73] Stef van den Elzen and Jarke J van Wijk. 2013. Small multiples, large singles: A new approach for visual data exploration. In *Computer Graphics Forum*, Vol. 32. 191–200. <https://doi.org/10.1111/cgf.12106>
- [74] Chris Weaver. 2004. Building highly-coordinated visualizations in improvise. In *IEEE Symposium on Information Visualization*. IEEE, 159–166. <https://doi.org/10.1109/INFVIS.2004.12>
- [75] Chris Weaver. 2009. Cross-filtered views for multidimensional visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 16, 2 (2009), 192–204. <https://doi.org/10.1109/TVCG.2009.94>
- [76] N. Wong, S. Carpendale, and S. Greenberg. 2003. Edgelens: an interactive method for managing edge congestion in graphs. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)*. 51–58. <https://doi.org/10.1109/INFVIS.2003.1249008>
- [77] Hao Wu, Maoyuan Sun, Peng Mi, Nikolaj Tatti, Chris North, and Naren Ramakrishnan. 2018. Interactive discovery of coordinated relationship chains with maximum entropy models. *ACM Transactions on Knowledge Discovery from Data* 12, 1 (2018), 1–34. <https://doi.org/10.1145/3047017>
- [78] Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. 2014. Opinionflow: Visual analysis of opinion diffusion on social media. *IEEE Trans. on Vis. and Computer Graphics* 20, 12 (2014), 1763–1772. <https://doi.org/10.1109/TVCG.2014.2346920>
- [79] Yanhong Wu, Naveen Pitipornvivat, Jian Zhao, Sixiao Yang, Guowei Huang, and Huamin Qu. 2015. egoslider: Visual analysis of egocentric network evolution. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 260–269. <https://doi.org/10.1109/TVCG.2015.2468151>
- [80] Hao Zhang, Maoyuan Sun, Danfeng Yao, and Chris North. 2015. Visualizing traffic causality for analyzing network anomalies. In *Proc. of Inter. Workshop on Security & Privacy Analytics*. 37–42. <https://doi.org/10.1145/2713579.2713583>
- [81] Jian Zhao, Chidansh Bhatt, Matthew Cooper, and David A Shamma. 2018. Flexible learning with semantic visual exploration and sequence-based recommendation of MOOC videos. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3173574.3173903>
- [82] Jian Zhao, Michael Glueck, Fanny Chevalier, Yanhong Wu, and Azam Khan. 2016. Egocentric analysis of dynamic networks with egolines. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 5003–5014. <https://doi.org/10.1145/2858036.2858488>
- [83] Jian Zhao, Zhicheng Liu, Mira Dontcheva, Aaron Hertzmann, and Alan Wilson. 2015. Matrixwave: Visual comparison of event sequence data. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 259–268. <https://doi.org/10.1145/2702123.2702419>
- [84] Jian Zhao, Maoyuan Sun, Francine Chen, and Patrick Chiu. 2017. BiDots: Visual exploration of weighted biclusters. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 195–204. <https://doi.org/10.1109/TVCG.2017.2744458>
- [85] Jian Zhao, Maoyuan Sun, Francine Chen, and Patrick Chiu. 2019. Missbin: Visual analysis of missing links in bipartite networks. In *IEEE Visualization Conference*. IEEE, 71–75. <https://doi.org/10.1109/VISUAL.2019.8933639>