

Active Reading of Visualizations

Jagoda Walny, Samuel Huron, Charles Perin, Tiffany Wun, Richard Pusch, and Sheelagh Carpendale

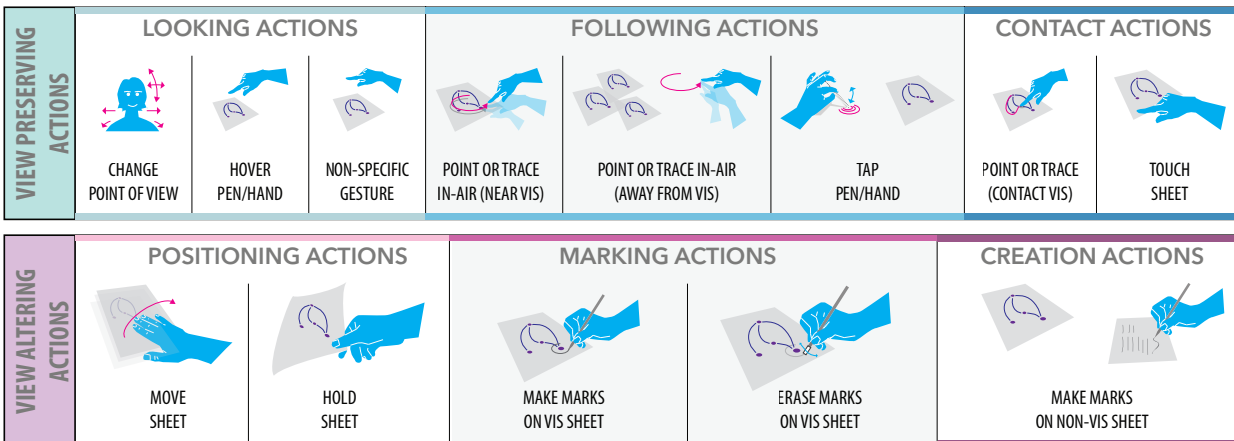


Fig. 1. Physical actions we observed during our study of active reading of node-link visualizations. Physical actions are ordered from left to right by increasing physical engagement.

Abstract—We investigate whether the notion of active reading for text might be usefully applied to visualizations. Through a qualitative study we explored whether people apply observable active reading techniques when reading paper-based node-link visualizations. Participants used a range of physical actions while reading, and from these we synthesized an initial set of active reading techniques for visualizations. To learn more about the potential impact such techniques may have on visualization reading, we implemented support for one type of physical action from our observations (making freeform marks) in an interactive node-link visualization. Results from our quantitative study of this implementation show that interactive support for active reading techniques can improve the accuracy of performing low-level visualization tasks. Together, our studies suggest that the active reading space is ripe for research exploration within visualization and can lead to new interactions that make for a more flexible and effective visualization reading experience.

Index Terms—active reading of visualizations, active reading, information visualization, spectrum of physical engagement

1 INTRODUCTION

We introduce the concept of *active reading of visualizations* using active reading of text as an inspiration. By studying whether the ideas of active reading of text might be applicable to the reading of visualizations, we open the door to the possible benefits of applying ideas from the extensive active reading research to challenges in reading visualizations, such as visualization literacy [12, 38, 40]. Starting with a qualitative study, we observed the actions people employed to read visualizations. We developed an activity spectrum, which relates to active reading of text, organized in Figure 1 as a spectrum of physical engagement. We then investigated possible benefits of supporting some of these actions for node-link graph visualization tasks. We found that people can answer more accurately when provided with support for these actions.

Visualizations are usually carefully designed to support specific purposes, such as acquiring information, understanding data-based stories, or gaining new insights about data. However, regardless of how well-crafted a visualization is, some effort is required to achieve the intended

aim in reading the visualization. Such effort can be challenging in both rudimentary and advanced reading. For instance, although our knowledge about creating perceptually accurate visualizations is growing [68], there are still concerns about visualization literacy [12] and challenges in supporting novices using visualizations [23, 40]. Supporting people in creative insight generation through visualization remains a significant, ongoing challenge [49, 59].

Readers of text face similar challenges. One way they cope with these challenges is through the established concept of *active reading*. This encapsulates the purposeful, engaged reading of text at a range of levels, from elementary reading for comprehension to advanced reading across multiple sources to generate new ideas [3]. Active reading is supported by various strategies that can take place internally in the mind or can be further aided by externalization, *i.e.*, the act of making one’s thoughts visible to support cognition [37]. An example of an internal active reading strategy is deciding to focus on specific parts of a text, such as verbs or key points in an argument. Examples of externalization-based active reading strategies are: highlighting parts of a text, making annotations within a text, or taking notes about the text. Active reading strategies are often discovered independently by readers but are also developed and taught in schools [45] and are often suggested as study skills for students (e.g. see [1]).

In this research, we propose and explore the concept of *active reading of visualizations* as a parallel to active reading of text: *active reading of visualizations is the purposeful, engaged reading of a visualization that combines internal and externalization-based reading strategies with available interactions to gain deeper comprehension.*

The first step of our exploration was to observe whether and how people actively read visualizations via a qualitative study of people reading paper-based node-link visualizations (study S1). We found

- Jagoda Walny, Tiffany Wun, Richard Pusch and Sheelagh Carpendale are with the University of Calgary. E-mail: {jkwalny, twwun, rapusch, sheelagh}@ucalgary.ca
- Samuel Huron is with Telecom ParisTech and the University of Calgary. E-mail: samuel.huron@telecom-paristech.fr
- Charles Perin is with City, University of London and the University of Calgary. E-mail: charles.perin@city.ac.uk

that people do read visualizations actively, and that they used a wide variety of actions (Figure 1). Combining these results with interview data, we gathered an initial set of visualization-specific active reading *goals* that group the low level actions under higher level goals. The second step was to better understand the potential benefits of supporting active reading of visualizations. We studied one group of actions that emerged from our first study, *marking and creation actions* (making and erasing marks), as these actions are involved in the higher-level goals of *decoding* and *analyzing* a visualization. To support these actions, we implemented a freehand annotation layer on top of a node-link visualization that participants decoded and analyzed in order to perform a set of low-level graph reading tasks (study S2). Results show that supporting marking and creation actions improved accuracy in performing our sample tasks with only a minor time cost, and that this effect is greater for larger, more difficult tasks. This indicates that supporting people in actively reading visualizations may facilitate their use of visualizations. In summary, we contribute:

- The introduction and definition of active reading of visualizations;
- Evidence that people use a variety of active reading actions to achieve various goals when reading visualizations (study S1);
- Evidence that there are benefits to supporting active reading goals in interactive visualizations (study S2).

We conclude by discussing our results in context of Bertin’s [9] stages of reading graphics and Adler’s [3] levels of active reading of text. Our studies suggest that some advantages attributed to active reading of text may apply to reading visualizations.

2 BACKGROUND: ACTIVE READING OF TEXT

We introduce the concept of active reading of visualizations by drawing on parallel concepts from active reading of text.

2.1 Active Reading of Text

Adler [3] defines active reading of text with a focus on internal mental activities, describing it as “the asking of questions” about a text. Thus active reading is the process of reading while being deliberately engaged with and thinking about the text. Questions are asked and answered differently depending on the reader’s goals, efforts, and skills. Adler identifies four levels at which a reader’s goals can differ. These levels capture the wide applicability of active reading from basic understanding of text to developing new ideas within entire subject areas.

Elementary level: Pertains to basic literacy about a text. Strategies such as circling all character names in a story and underlining new vocabulary [55] are taught in reading education [45].

Inspectional level: Involves gaining a picture of the text using, for example, systematic skimming strategies to understand the structure and type of a book. One might examine the book cover, title, subtitles, figures, and genre of text [55].

Analytical level: Has the goal of increasing one’s understanding of the text, for example by “asking many, and organized, questions” while reading [3]. This is related to a technique called *close reading*, which is a systematic way of directly reading a text “to uncover layers of meaning that lead to deep comprehension” [13].

Syntopical level: Involves reading multiple sources and constructing new analyses that might not be present in any individual book [3], for example in knowledge work when multiple documents must be cross-referenced [2]. This can be linked to *distant reading* [46], a technique in the digital humanities that uses statistics and visualizations to understand texts by examining their features and structures [33].

Deep engagement with the text is what Pearson et al. [51] identify as the *primary task* of active reading. However, active reading is often supported by other activities, which Pearson et al. call *secondary tasks* [51]. These secondary tasks often take the form of externalizations such as note-taking, annotation, or marking up the text [2, 47, 50]. They are also used for teaching reading comprehension [55].

2.2 Digital Support for Active Reading of Text

On paper, secondary tasks for active reading can be considered “lightweight interaction” [51] because the reader does not need to consciously think about them. However, in digital environments, support-

ing these tasks is more challenging. For instance, highlighting a word in a digital reader typically requires finding and pressing a button to enter a highlighting mode before selecting a cursor. This contrasts with the fluid motion of highlighting with a physical highlighter. Because of this, numerous projects have investigated the creation of digital active reading environments, including: XLibris [57], LiquidText [64], GatherReader [27], Matulic and Norrie’s pen-and-touch active reading environment [42], and systems that support close reading [33]. The design of some of these digital text reading environments was informed by studies of active reading tasks [29, 63]. Some have also suggested leveraging the digital environment to support new tasks, such as Text-Tearing to create space for annotations [73] or word-scale visualizations to allow cross-referencing within a text [22].

Whereas the main challenges in digitally supporting active reading of text stem from matching the accustomed fluidity of reading on paper, information visualizations have a smaller paper-based legacy, having gained much of their popularity in their digital form. Information visualizations also do not have a standardized structure like text does; however, they do share some parallels with text. Visualization literacy [12] is a term borrowed from text. Collections of text documents are a common type of visualization dataset, e.g. for sensemaking [7], for the digital humanities [28], or for personal use [70]. The content of text can be visualized, as in, e.g., the content of documents in DocuBurst [16], the content of a novel in Writing Without Words [56], and the sonic topology of poems in Poemage [43].

2.3 Relating Active Reading to Visualization Research

The term “active reading” has not previously been associated with visualizations. However, some interactions in visualizations, such as annotation mechanisms, can be discussed in terms of active reading. In visualization, such mechanisms have primarily been designed with the intent of communicating thoughts to others during collaborative [41] and crowdsourced analysis [24, 69], public discussion of narrative visualizations [30], and for authors of narrative visualizations to explain data [60]. Sense.us [24] is an exception, in that it went beyond offering the possibility of placing of text in sidebars or behind icons by providing a freeform graphical annotation layer. Interestingly, although this layer was intended for collaboration, some people seemed to prefer this graphical overlay for personal use [24]. Our exploration of active reading may offer an explanation for this personal use: freeform annotation provides rich support for active reading techniques.

While we recognize the considerable benefits gained from active reading of text, our research goal is to unveil active reading behaviours as they occur when people are reading visualizations. Thus, rather than simply apply behaviours from active reading of text directly to visualizations, we study whether and how people actively read visualizations.

3 METHODOLOGY AND RATIONALE FOR THE STUDIES

Here we briefly explain our methodologies and the relationship between them. McGrath [44] states that generalizability, precision and realism are all desirable characteristics of empirical work, but are impossible to achieve with a single study; to arrive at a complete understanding, multiple studies are needed. Following this, we took the approach to conduct two linked studies with different purposes.

Our first step towards *laying the groundwork* of active reading of visualizations is to establish whether people naturally read visualizations actively. To achieve *realism* rather than generalizability, qualitative studies with small numbers of participants are increasingly valued [28, 48]. Observing people using non-digital artifacts has been a fruitful method for expanding our perspectives on the possibilities of interactive systems, including visualizations (e.g., [14, 31, 32, 67]). Following this approach, we designed a qualitative study with 6 participants, study S1, to explore *if* and *how* people actively read visualizations.

Once we laid out the space of active reading actions and goals via study S1, we designed study S2 to take a step towards *precision*. Because moving towards precision narrows the scope that a single study can encompass, we chose one subset of actions to study: marking and creation. In S1, these corresponded to the active reading goals of *decoding* and *analyzing* visualizations, which are of particular interest




Fig. 2. The table set up: piles of paper visualization sheets (bottom); four transparencies (middle); tools (top) are available for participant use.

to visualization researchers. One of the best ways to progress towards precision is through controlled lab studies in which participants are asked to perform low-level tasks (e.g., [15, 18, 61]). Following this approach, with S2 we evaluated the accuracy and completion time of 18 participants for two low-level graph reading tasks, with and without active reading support for marking and creation actions.

In choosing one of the many possible visualizations for this pair of studies, we wanted a visualization that: 1) is at least moderately common; 2) is easy to explain; 3) can be used to represent a dataset that is familiar to people in general; 4) is feasible for use in both study S1 and S2; and 5) is commonly enough studied in the visualization community that tasks of varying difficulty are available in the literature. We used node-link diagrams portraying social networks where the nodes represent people and the links represent the relationships between them, because they fill all of these conditions. This leaves room to study additional types of visualizations in future work in this space.

4 S1: OBSERVING ACTIVE READING IN VISUALIZATION

We designed our qualitative study S1 to explore *if* and *how* people actively read visualizations. To avoid the constraints and advantages of software, we observed people working on paper-based visualizations.

4.1 S1: Participants, Set-up and Materials

We recruited six participants (3F, 3M, 18-24 years of age) using campus-wide posters and word of mouth. We set up a table in a quiet room. Two video cameras recorded the table, providing a top and a side view.

We used a social network dataset with ten nodes (persons) and four types of links (knows, likes, loves, and dislikes). We gave participants four distinct node-link views of this network, one for each type of link (see Figure 3). Each view was available on letter-size paper and on transparencies that could be layered to view several types of links at the same time. We provided participants with a set of materials and tools: blank paper, blank transparencies, tracing paper, water-soluble markers, water and a cloth for erasing marks on transparencies, a regular pen, erasable felt pens, scissors, tape, post-it notes, and paper clips. We covered the table in white paper to make the transparencies easy to see.

4.2 S1: Procedure

First, participants filled out a consent form and a demographic questionnaire. Then the experimenter explained the visualization, taking care to make a single mark on one of the transparencies to make it clear to the participant that the materials were not too precious to be written on. Next, the three phases of the experiment were: (1) questions, (2) problem solving, and (3) interview. To start phase 1, the experimenter told participants that they could use any tools or materials on the table, and then gave them a sheet containing five questions related to the visualization and space to write their answers. The questions asked participants to find the most liked and most disliked people, those who know the fewest people, those who love someone who doesn't love them back, and asked if there are any close-knit groups of friends.

These five questions were refined during pilots to require some counting, comparisons, decoding and interpretation of the visual mapping.

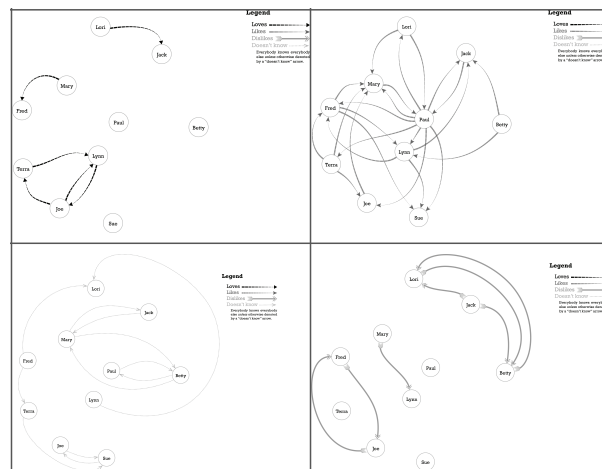


Fig. 3. The four views (showing different link types) provided to participants on separate letter-sized sheets of paper.

They also covered a variety of common low-level visualization tasks such as finding entities, outliers, clusters, and highly connected and poorly connected entities. At the end of this first phase, participants were asked to explain their strategies for answering the questions.

We designed phase 2 to observe participant behaviour in a more difficult problem solving situation. During this phase, participants used the same materials to solve the problem of how to organize a social evening with all of the people depicted in the social network with “minimal drama” due to incompatible relationships. Once they were done, participants were asked to explain their problem solving strategy.

In phase 3, participants filled out a questionnaire about their active reading habits when the information was on paper, on mobile devices, and on personal computers. This questionnaire was given last because the questions suggested active reading behaviours and we did not want to bias the results by asking them think about active reading beforehand. This questionnaire formed the basis for a semi-structured interview about the participants’ active reading habits and how these habits related to the strategies they used to read the visualizations.

4.3 S1: Analysis Method

Two researchers independently open-coded [62] phase 1 and 2 of the video data for two different participants, looking at the synchronized top and side recordings simultaneously. These researchers then iteratively coded additional participant data and discussed their codes until reaching consensus. At this point, one researcher coded the remaining videos, discussing questions as they arose to maintain consensus. The interview in phase 3 was then analysed based on these codes.

4.4 S1: Results: Physical Engagement Spectrum

All participants demonstrated considerable physical movement when reading the visualizations. As participants became more physically engaged with the artifacts on the table, they exerted greater control over what was in their field of view. We therefore split our codes into two groups: *view-preserving* and *view-altering* actions. Figure 4 shows the counts of actions for each participant for phase 1 and 2 of the study. Ordering these actions according to increasing physical engagement results in the *physical engagement spectrum* shown in Figure 4.

4.4.1 View-Preserving Actions

View-preserving actions maintain the arrangement and content of the visualization sheets on the table. Through these actions participants used their body to temporarily augment their view, e.g. by tracing virtual links between nodes with their fingers to follow relationships.

Looking actions are those associated with participants exerting the least control over the visualization: looking at, but not interacting with, the visualization. We coded for moments when participants *changed their point of view*, usually through head turns or posture changes. We also coded for moments when gestures accompanied




Fig. 4. Physical engagement spectrum with count of actions by phase for each participant. The coloured stripes on the right-hand side indicate the higher-level goals participants had when performing these lower-level reading actions, based on the interview data.

viewing only. These included *hover* gestures, in which participants held their hands in a specific, non-moving position; and *non-specific gestures*, in which participants made hand gestures that did not directly reference a particular aspect of the visualization.

Following actions are those in which participants performed gestures that clearly referenced particular aspects of the visualizations, without any contact. Participants tapped their pen or hand in the air. They pointed to nodes or traced relationships within the visualized network both *away* from the visualization and *directly above* the visualization, where it was clear which nodes or edges were being referenced.

Contact actions are those in which participants made contact with the visualization, either with their hand or mediated by a tool such as a pen. We observed more following actions, such as *tracing a relationship* while contacting the sheet, and *touching the visualization* — either touching the marks or the sheet. Some participants touched a sheet with one hand while looking at another sheet or touched one mark with one hand while tracing relationships with the other hand.

4.4.2 View-Altering Actions

View-altering actions are those that changed the participants’ view in a semi-permanent or permanent way. This refers to changes to the spatial arrangement of sheets and to the markings visible on the sheets.

Positioning actions include those in which participants held a sheet temporarily in a position. Participants either *moved sheets* around or




Fig. 5. Examples of marks made (a) on top of the visualization and (b) away from the visualization.

held the sheet up above the table. They moved sheets into and out of their field of view, closer and further away, and they rotated, translated, tilted, and layered the transparent sheets.

Marking actions are those in which participants marked the visualization sheet, *making* or *erasing* marks on the visualization. This includes modifying visual encodings, annotating, tracing over nodes or edges with a pen, and erasing or scribbling out previously made marks.

Creation actions are those in which participants *made* new representations (see Figure 5), usually away from the visualization on a blank sheet of paper. These include sketches, lists of node names in various arrangements, textual phrases, and full textual sentences.

4.5 S1: Interpretation

The low level actions listed in Figure 4 demonstrate that there are people who, when asked to use a visualization to answer some data questions, make use of a variety of physical actions to answer these questions. All participants used a variety of actions from the physical engagement spectrum to actively support their reading during problem solving. Figure 6 shows how the activity sequences of these low level observable actions varied not just from person to person but also from one problem to the next. By providing participants with a paper-based environment, we gave them the freedom to support their reading actively in ways that might not be available digitally.

Participants performed actions with increasing physical engagement as questions got more difficult. In phase 1, question 5 was the most difficult and time-consuming question, requiring synthesis of multiple pieces of information and making decisions. For this question we saw increased physical engagement in comparison to simpler questions, for example more tracing and addition of marks onto the visualization. For phase 2, where participants had to come up with a creative solution to a problem, we observed different behaviours. They tended to first arrange the visualizations within their field of view, and then use a separate sheet of paper to create externalizations that represented the information and their thoughts to support their problem solving strategy.

4.5.1 Goals for Active Reading of Visualizations

As we could not observe participants’ internal mental processes, we combined our observations of physical actions with interview explanations from our participants to infer some of the reasoning underlying the actions we observed. Similarly to the variation shown in Figure 6, participants might use one action for one goal (i.e. purpose) for a given problem and apply the same action for a different goal at another time. The actions were used as possible processes which might be useful in different situations and for different goals. These goals include: recognizing, tracking, re-organizing, decoding, and analyzing.

Recognizing: All participants spent considerable time simply looking at the visualizations. The intent appeared to be to discover what could be understood with minimal engagement. In coding for observable actions, we coded simple looking activities, when they involved




Fig. 6. Sequence of participant actions, colour-coded by category in the depth of physical engagement spectrum in Figure 4. The width of each coloured block represents the duration of one action or of successive similar actions. This demonstrates how varied participants' actions were, not only for different participants but even for different tasks for the same participant.

head movement, as *changing point of view* actions. Simply looking at the visualization, with or without head actions, accompanied all low level actions except creation actions where the participants were looking at their new creations instead. Even when limited to looking with head movement, note the prevalence of this activity in the top row of Figure 4 and its high frequency in Figure 6. The view-preserving action of looking with head movements to change point of view contributes both to recognizing and to reorganizing (below) because participants were using this action to change their point of view.

Tracking: Keeping track of the elements in a visualization is challenging, particularly when there are many similar-looking elements. It requires focused attention and the use of memory. To track, participants used different tracing and following techniques. P2 noted, “*I find that with this [the vis] I need to highlight more and do more annotation [...] because it doesn't, it's kind of confusing to me to just look at.*” Many of the actions in the physical engagement spectrum contributed to the different activities that participants used to keep track of their understanding of the visualization. These included: *view-preserving actions* such as *hovering with either pen or hand, pointing or tracing in air (either near or far), tapping, and pointing or tracing on the vis*; and some of the *view-altering actions* including *making marks on vis sheet, erasing marks on vis sheet, and making marks on other sheets*. Tracking was a frequent and diverse goal.

Participants also kept track of which information had been visited. They used physical actions like annotating or highlighting nodes or links as they were visited. Another tracking technique is to save the location of an object of attention. This way of offloading memory was achieved through pointing, touching, tapping and sometimes involved bimanual actions. For example, P2 held a node with one finger while pointing to its connected nodes as she followed outgoing links. These actions are similar to those observed in the context of physical visualizations [34].

Reorganizing: Searching for and relating information between multiple views can be difficult. To cope, most participants spatially reorganized the documents to set up the reading environment that suited their personal preferences and needs. We observed all participants arranging views by moving the individual sheets primarily using *positioning actions, move sheet and hold sheet* from the physical engagement spectrum. Figure 4 shows that all participants moved sheets for both phase 1 and 2, ranging from 3 to 23 times. All participants preferred their own unique arrangements. Arranging views placed a manageable subset of views in front of a reader and may also place views within spatial memory to reduce the time to search for views. Some participants mentioned wanting greater reorganizational freedom. For instance, P2 wished to move nodes instead of redrawing them: “*If I could have just actually moved them [the nodes] and put them together myself [...]*”.

Decoding: Creating a visualization involves encoding data and developing a mapping from data to visual, spatial entities. When reading a visualization the inverse is true. Part of the challenge is to decode the visualization so that the visual-to-data mapping is understood. This can be challenging when the encoding is unfamiliar. One technique participants used was to re-encode the information to better suit their own needs or internal representation. Most participants re-encoded some information. P2 related this technique to her own active reading practices, which favour note-taking: “*I like to put it in my own order, I guess? To make it easier for myself to understand.*” The actions that were used for the goal of *decoding* typically included *view-altering*

actions, especially making marks on vis sheet, erasing marks on vis sheet, and making marks on other sheets. P2 added a new encoding by superimposing the “dislikes” relationships onto the “likes” sheet and drawing coloured crosses, a different encoding than the one used in the original visualization. Using re-encoding techniques to help decode the visualization may be a response to the difficulty of adapting one encoding for multiple types of problems. This is in line with the Congruence Principle for effective graphics: that “the structure and content of the external representation should correspond to the desired structure and content of the internal representation” [65].

Analyzing: Whereas the previous active reading goals support reading information that is already present in the visualization, deeper reading can involve integrating or synthesizing new information. This goal of analyzing usually involved *making marks on other sheets* and sometimes included *making marks on vis sheet*. In Phase 1, P5 annotated each node as he counted its links. In Phase 2, P3 wrote a list of all of the node names on her free sheet and crossed them off when she had addressed them, saying, “*That was so I don't have to keep it all in my head.*” Analyzing often involved recording new information. Reading a visualization may require making interim inferences, resulting in many thoughts about the information and how it relates to the task at hand. Storing this information in working memory can be overwhelming. To tackle this, participants counted nodes and links or made calculations, then recorded counts or calculations by making marks. Participants used blank sheets to record partial or full solutions. These solutions usually referenced elements of the visualization, such as specific nodes or relationships. Participants sometimes used the blank sheet simply to offload thoughts about the visualization or problem they were solving.

4.5.2 S1: Limitations

Through this study, we observed people in the context of engaging with node-link visualizations, and using and adapting their own active techniques as necessary to the tasks they were given. We observed some comparable activities to active reading of text (marking, highlighting, making notes, etc.). However, since the representations of visualizations and text are so different, more research is needed to better understand the extent to which existing knowledge about active reading of text can apply to active reading of a range of visualizations and tasks.

4.5.3 From Qualitative Results to a Quantitative Study

In this qualitative study, we observed that within the context of visualizations, people had their own active strategies that they used to help themselves solve the assigned tasks. These were purposeful, engaged actions that helped participants understand and work with the visualizations by combining active strategies (as observed via physical actions) with internal strategies (as revealed via interviews).

Given the evidence from S1 that people use active techniques when reading visualizations on paper, the next step is to determine whether it would be useful to support such techniques with digital visualizations. Any of these active reading techniques could be studied. To choose which technique to study first, we focused on actions used during Phase 1, when participants were exposed to the visualization for the first time and had to decode the visual mapping and develop their initial understanding of the data. We also considered that view-preserving actions could be used on existing digital visualizations without augmentation. However, some view-preserving actions have corresponding

view-altering actions. Specifically, tracing above, near, and on the surface can be grouped with marking and creation actions; in combination, this group of actions had widespread use (see Figure 4).

For our next step, we studied support for freeform annotation because providing the ability to make marks on or adjacent to a visualization directly supports marking and creation actions. Freeform annotation also makes it possible to employ a variety of other active reading techniques that we observed: tracking, decoding, reinforcing encodings, adding a new encoding, transforming and re-encoding subsets of information, as well as analyzing actions such as recording counts or calculations, recording partial or full solutions, and off-loading thoughts. Another reason for choosing annotation is that this type of action can be supported with simple, familiar interactions with an easy learning curve.

5 S2: STUDYING BENEFITS OF ACTIVE READING SUPPORT

We conducted a within-participants full factorial design quantitative study (S2) to determine whether providing freeform annotation to support active reading can improve speed or accuracy when performing low-level visualization tasks. We chose to study a basic active reading support technique, providing a freeform annotation layer on top of an interactive visualization, because it closely replicates what can be done on paper (paralleling study S1), it supports both marking and creation (from study S1 results), and it relates to the graphical overlay people appropriated for personal annotations in Sense.us [24].

5.1 S2: Factors

The experiment included three factors: *CONDITION* was either *BASELINE* (standard graph visualization interface with touch interaction) or *ACTIVE* (standard interface augmented with an annotation layer with pen and touch interaction). *TASK* was either *DEGREE* (node-counting task) or *REACH* (accessibility task). *N* (N20, N40, N80) was the number of nodes in the graph, i.e. the visualization complexity.

The basis of our experiment was a set of randomly generated static, undirected, unweighted node-link graphs with curved edges, visualized using D3 [10]. To ensure consistency across trials, we generated a set of graphs with constraints (using D3’s force-directed layout) that all participants used in randomized order within each task. We implemented two basic touch interactions to aid readers in following connections where nodes or edges are placed closely together and where edges cross: touching a node highlighted its connected nodes and the edges between them, and dragging a finger over an edge highlighted that edge and the nodes connected to it (see Figure 7).

5.1.1 S2 Factor: Conditions

The *BASELINE* condition consisted of this basic touch-enabled implementation. The *ACTIVE* condition was identical but with the addition of a layer on which marks could be made using a pen input device. We call this layer the *freeform annotation overlay*. We implemented the freeform annotation overlay illustrated in Figure 8 as an SVG group placed on top of the D3 visualization. Participants could draw freehand SVG paths over the graph in three different semi-transparent colours (yellow, pink, and blue) and three different thicknesses, and could also erase paths. Participants used a button-based palette at the right hand side of the visualization to switch pen properties. Pen properties persisted across study trials to allow for personal preferences.

5.1.2 S2 Factor: Tasks

In visualization it is common to study low-level tasks that can be combined to accomplish more complex visualization-reading operations [5]. We selected two tasks from Lee et al.’s taxonomy of graph tasks [39] that i) both involve counting and could be answered using the same input modality (providing a number); ii) are of different categories in the taxonomy [39]; and iii) are among the most discriminating tasks that involve counting [21]. The two tasks are *DEGREE* (the number of nodes that have the maximum degree) and *REACH* (the number of nodes that can be reached from the selected node in two or fewer steps).

DEGREE is a topology-based counting task about adjacency (connectivity) of nodes [39]. This is a compound task that requires participants to both *find the maximum degree* and *count the number of nodes with*




Fig. 7. Illustration of the two kinds of touch interaction included in both *BASELINE* and *ACTIVE* conditions. Nodes and edges remain highlighted until either another element or the background is touched.



Fig. 8. Using touch interaction to highlight links while highlighting on top of the graph using the freeform annotation overlay.

that degree. With this task, possible causes of errors are losing track of which nodes have been counted and miscounting edges. Adjacency tasks that consist of finding the most connected node have been featured in several previous studies [21, 25, 54, 58], as have variants on this task [4, 20], including our *DEGREE* task [26].

REACH is a topology-based counting task related to accessibility [39]. With this task, a possible cause of errors is counting nodes more than once. Accessibility tasks appear in numerous studies, often in the form of finding a path between two nodes [20, 21, 25, 26, 54, 58, 72].

These two tasks have the potential to be helped by active reading support because we found in S1 that people *keep track* of nodes and edges, often by making marks on top of the visualizations. Also, these two tasks often reveal differences between visualization techniques when the number of nodes increases [4, 26, 58].

5.1.3 S2 Factor: Numbers of Nodes

We tested graph sizes of N20, N40 and N80 (doubling the number of nodes each time). This was based on our observations in S1 that people used different active reading strategies when task difficulty increased, suggesting that the benefits of *ACTIVE* could depend on the difficulty of the task. Varying the graph sizes ensured that we provided tasks of varying difficulty levels, and is in line with similar studies [21, 26, 58, 72]. As is often the case (e.g. see [53]), we did not balance *N* and participants were presented with datasets of increasing complexity (N20, then N40, then N80).

5.2 S2: Hypotheses

Our hypotheses for the experiment were as follows:

- H1 ACTIVE will result in lower error rates for both tasks.** We expect that having access to the freeform annotation overlay will help participants keep track of the nodes they have counted, resulting in greater accuracy.
- H2 ACTIVE will have greater impact with larger graphs for both tasks.** Based on the results of S1, we expect that the freeform annotation overlay is more useful when there are larger numbers of nodes because of greater demands on working memory.
- H3 ACTIVE will be slower than BASELINE for graphs of larger sizes and will be only minimally slower for graphs of smaller sizes.** We expect that drawing on the freeform annotation overlay will cost time in accordance with the active reading strategy used.

5.3 S2: Design, Setup, and Participants

The experiment consisted of two CONDITION blocks. Each block was split into two TASK blocks, which were themselves split into three N blocks. The order of CONDITION \times TASK was randomized across participants, and N appeared in increasing order within each block. Each CONDITION \times TASK \times N block consisted of 4 recorded repetitions. In total, the experiment consisted of: 16 participants \times 2 CONDITION (BASELINE, ACTIVE) \times 2 TASK (DEGREE, REACH) \times 3 N (N20, N40, N80) \times 4 repetitions = 768 measured trials.

Our dependent variables were *time* and *error*. Time is the time spent to perform a measured trial. Error magnitude is a percentage of the correct answer: $error = 100 \times \left| \frac{answer - correct}{correct} \right|$. This error measure has been used for similar graph-related tasks [20] as it provides information about the relative magnitude of the error rather than binary correctness. This is quite useful for tasks where *correct* is a relatively large number.

We recruited 16 student participants not involved in S1 (aged 18 – 35 years, 7 females) via posters displayed on a university campus, word of mouth, and mailing lists. Participants sat in a quiet room at approximately 30cm from a 24 inch Wacom Cintiq 24HDT display with pen and touch using the Chrome browser full screen at 1920x1200 resolution (see Figure 9). The experiment was video recorded. The whole experiment took approximately 90 minutes. Participants received a \$20 remuneration for their participation.



Fig. 9. The physical study setup used a 24 inch pen-and-touch display. A camera captured footage of participants' hands on the display.

5.4 S2: Procedure

1. Preamble. The experimenter read an introductory script explaining the experiment. Participants then filled out a short demographic questionnaire. Participants proceeded through the study on their own, following written instructions on the screen and using the touchscreen to progress. An experimenter was available for questions at all times.

2. Tasks. When starting a new CONDITION block, participants had unlimited time to familiarize themselves with CONDITION using a test graph. In the ACTIVE condition, the interface prompted participants to make use of the pen. Each CONDITION had two TASK blocks. When starting a new TASK (twice per condition), an instruction screen explained the task. Participants then performed two training trials and they were encouraged to ask questions at this time. Once ready, they proceeded to the block of measured N20 trials.

We measured the time to complete a trial from when the participants pressed a button labelled 'GO', which made the next graph appear, to when they input and confirmed their answer. To ensure that participants understood the task, the interface showed correct answers after each trial. Once the N20 trials completed, participants were asked "How easy did you find this task?" and answered using a Likert scale (1: very easy, 5: very difficult). This structure was repeated for N40 and N80.

3. Epilogue. After completing all trials, participants completed a questionnaire about their experience and their active reading habits.




Fig. 10. Error and time 95% confidence intervals for both conditions for both tasks and by number of nodes. Black dots indicate the best estimate while confidence intervals convey effect sizes. The gray areas show error and time pairwise mean comparisons, that is, the participants' differences between the two conditions.

5.5 S2: Results and Analysis

To report the results of our study, we follow the recommendation from APA [6] and base our analyses on *estimation* using bootstrapped [36] confidence intervals [17] instead of p-values. A 95% confidence interval contains the true mean 95% of the time and conveys effect sizes [17], making it possible to *estimate* differences. This approach has been recommended for reporting statistical results in HCI over the traditional null hypothesis significance testing (with p-values only), which leads to dichotomous thinking [19]. It has seen increased use in HCI and visualization (e.g. [11, 18, 35, 66, 71, 74]). We pre-specified all analyses before conducting the experiment and tested on pilot data. This includes the R scripts used for parsing the data, computing confidence intervals of pairwise mean comparisons, and generating drafts of figures.

Figure 10 shows error magnitude and completion time by CONDITION, TASK, and N, along with within-participant mean differences between BASELINE and ACTIVE. The best estimates for each participant were computed by using the mean of the participant's four measured trials within each CONDITION \times TASK \times N block. Black dots are mean point estimates, i.e. best guesses, and black lines represent confidence intervals, whose width conveys effect sizes. For readers familiar with p-values, a pairwise mean comparison confidence interval that does not cross the 0 vertical line can roughly be interpreted as $p < .05$. We qualify these results as *strong* if the pairwise comparison shows a confident effect, that is, the confidence interval does not cross the vertical axis at 0% (error magnitude) or 1.0 ratio (time). We qualify the results as *weak* if there is probably an effect, but the confidence interval is wide or crosses the 0% (for error) or 1.0 (time) vertical axis.

5.5.1 Analysis

Overall, the ACTIVE condition reduced the magnitude of error for larger graphs with a time cost of up to 2.0 \times . The discussion that follows is based on the pairwise mean comparisons shown in Figure 10.

For DEGREE, ACTIVE resulted in lower error magnitude than BASELINE for all N. The effect was strongest for N80, with a 12% reduction in error magnitude; small ($< 10\%$) and weak for N40; and small ($< 10\%$) but strong for N20. This confirms our hypotheses H1 and H2 for DEGREE. The reduction in error magnitude came with an increase in time (strong effect of CONDITION on time for all N). For N20, there was a small ($\approx 1.25\times$) cost of time to complete the task. For N40 and N80, the time cost was close to $2.0\times$. This partially confirms our hypothesis H3 that there would be a larger time cost for larger graphs, however, it differs from our original hypothesis that there would be a minimal time cost for the small graph.

These results indicate that participants were able to use the freeform annotation overlay to improve their accuracy in finding all nodes of the maximum degree in a graph. We expected the difficulty of DEGREE to increase with the complexity of the graph. Video records confirm that participants leveraged the freeform annotation overlay to more accurately keep track of previously visited nodes, of the degree of visited maximum-degree nodes, and of the running total of maximum-degree nodes. The longer times spent with ACTIVE make sense because of the time it takes to draw annotations.

For REACH, ACTIVE resulted in strong reduction in error magnitude for N80 ($\approx 5\%$). For N20 and N40 the difference was weak and $< 2.5\%$. This confirms H2: the effect of ACTIVE is greater for larger graphs; however we cannot confidently confirm H1 that ACTIVE has an effect at all graph sizes. We observed a strong time cost of $\approx 1.5\times$ for ACTIVE for N20, however, we observed no differences in time for N40 or N80. This does not confirm our hypothesis H3 that ACTIVE would be slower than BASELINE for all N.

These results indicate that while N80 was difficult ($> 10\%$ error in the BASELINE condition), participants leveraged the freeform annotation overlay to find a more accurate answer. We expected the difficulty of REACH to increase with the complexity of the graph. Video records indicate that the freeform annotation overlay allowed participants to mitigate this difficulty when the number of nodes to count was large.

For both DEGREE and REACH, the error for N40 appeared to be larger than for N80 with ACTIVE. This may be because we did not counter-balance N: N20 was simple enough that no active reading strategy was needed, but N40 and N80 were more difficult. Once participants reached N40, they needed to find an effective and comfortable active reading strategy. By the time they reached N80, this strategy would have been well established as compared to the beginning of N40.

Overall, these results show that people can leverage a freeform annotation layer to read visualizations more accurately, though at some time cost. This is particularly useful at higher graph complexities where it appears that relying on internal mental representations and memory becomes more difficult. Because participants were not trained in using any specific active reading support techniques, this study shows the ability of people to spontaneously create their own support techniques as needed when the opportunity is provided.

5.5.2 S2: Limitations

Based on the results of S1 and on the rich variety of active reading techniques for text, active reading appears to be a highly varied, individualized activity that depends on a reader's goals and skills. For this reason, we did not impose any particular active reading techniques on participants. While this allowed us to study whether the availability of a freeform annotation overlay has an effect on reading visualizations, it also created variance in the techniques used by participants. Therefore, some participants may have used active reading techniques that were not optimal for the task at hand. For instance, we found that one participant highlighted every counted node and link in the DEGREE task, greatly increasing completion time.

6 DISCUSSION

We investigated if the concept of active reading is applicable in a visualization context, what this concept might mean in a visualization context (S1), and if it provides measurable benefits for reading visualizations (S2). We discuss these results in the context of Bertin's [9] stages of reading graphics and Adler's [3] levels of active reading of text.

6.1 Using External Actions to Read Visualizations

The results of S1 demonstrate that there are active reading actions that some people do apply in some visualization contexts. The paper-based setup combined with the permission to make marks allowed for considerable freedom in re-positioning views, in using transparencies for filtering, and in annotating. In S1 all participants, spontaneously and unprompted, manipulated and marked the views to varying degrees and in a variety of ways. This observation of spontaneous actions that parallel active reading actions is important because it suggests that at least some portion of visualization reading can be aided by external actions and that some people, when given the opportunity, naturally use external actions to help themselves read visualizations.

The use of active reading does tend to lengthen the reading process. This is in part because of the additional actions involved, as in S2, where participants completed tasks more accurately, but also more slowly, with the freeform overlay. In visualization, increased time cost is often considered a downside. However, in active reading of text, speed is a less important factor, as more emphasis is placed on increasing comprehension, engagement, and insights into a text [3]. It is in these more complex goals that active reading may have a place in visualization as well.

6.2 Setting Active Reading of Visualizations in Context

From S1 we learned about the low-level observable physical actions (Section 4.4), and, combining these with interviews, extracted the higher-level goals (recognizing, tracking, reorganizing, decoding, analyzing) described in Section 4.5.1. These results relate to and extend Bertin's reading stages of graphics and Adler's levels of reading of text.

Bertin [9] describes a 3-stage process for reading paper-based data graphics: (stage 1) external identification of factors in the graphic that relate to the reader's existing knowledge; (stage 2) internal identification, in which the reader examines and assigns meaning to the visual variables' encodings of the data; and (stage 3) perception of pertinent correspondences, in which the reader performs basic data queries: elementary (single data item), intermediate (group of data items), and overall (overview of data in the graphic). Bertin specifically describes each of these stages as internal mental processes. Despite the added capacity for interactivity in information visualizations, much visualization work has followed this assumption that internal mental processes alone are sufficient to perform these reading tasks. Bertin's reading stages consider how to read, and his reading levels focus on questions asked in terms of data. The inspiration from active reading of text is that there are advantages to be gained by explicitly externalizing factors that contribute to these reading processes (as opposed to internal processes only). Our five high-level reading goals provide a framework for the lower level actions that we observed as people used externalization to help themselves read the visualizations.

Adler [3] identifies four levels at which an active text reader's intentions can differ (elementary, inspectional, analytical, syntopical — see Section 2). Adler's levels express increasing complexity and identify specific contexts in which to apply active reading techniques.

Figure 11 shows correspondences between our five goals for active reading of visualizations, Adler's four levels of active reading of text, and Bertin's three stage reading process for data graphics. Adler's levels contain recommendations for externalizations along with internal processes, while Bertin's stages are described as internal mental activities. However, the goals of the 3 stages and the first 3 levels are quite parallel. For example, Bertin's stage 1, external identification, has the goal of assessing what one can understand from one's prior knowledge. Adler's level 1, elementary, is also about assessing what one knows at the start of reading and adds that external actions, such as circling known characters in a story, can help establish what is known at the start. In our goals, Bertin's stage 1 relates to recognizing and Adler's level 1 relates to both recognizing and tracking. Bertin's stage 2 advises examining the encodings within the graphic (relating to recognizing, tracking with your eyes, and decoding), while Adler's level 2 suggests skimming to develop overview (a different behaviour) that also relates to recognizing, broad-sweep tracking, and decoding. Bertin's stage 3 involves reading and analyzing local, small group and global data




Fig. 11. Relating our five goals for active reading of visualizations with Bertin's [9] three steps of the reading process, and Adler's [3] four levels of active reading of text.

(involving tracking, decoding and analyzing). Adler's level 3 suggests questioning actions (also involving tracking, decoding and analyzing).

Note that a correspondence to re-organizing, which needs interactivity, is not present in either Bertin's or Adler's list, possibly because neither printed graphics nor text can readily be re-organized. Re-organization is available through interactivity in most visualizations and to some degree was simulated in S1 through separate views, transparencies and annotation. While Bertin did not mention re-organization in his stages of reading, he worked with the power of re-organization on comprehension by developing methods [8, 52] for reordering matrices. Bertin was fully aware of the role of graphic manipulation in the reading process, writing: "[Manipulation] is fundamental. It is the internal mobility of the image that characterizes the modern graphic", and that a graphic "is not drawn once and for all; it is constructed and reconstructed until it reveals all the relationships constituted by the interplay of the data." [8] Another exception is Adler's syntopical level, in which he refers to reading across many different texts for meta-analytic purposes. It would be worth investigating whether this parallel holds to situations in visual analytics where analysts work with many visualizations in multiple and coordinated views.

While Bertin's stages, Adler's levels, and our spectrum of physical actions (Section 4.4) are fundamentally different in terms of actions, Figure 11 shows that they can be related in terms of reading goals. Similar goals being achieved by a diversity of actions suggests that a great variety of actions, both internal (Bertin) and external (Adler, S1) can lead to successful reading and comprehension. We hope that this discussion will interest others in exploring which actions are best suited to which visualizations for different reading comprehension goals.

7 IMPLICATIONS FOR DESIGN

Our two studies, S1 and S2, suggest several interesting design opportunities. One of these is exploring how to best support active reading in digital environments. Some popular visualization interactions have parallels in active reading practices (i.e. brushing and linking, highlighting). A possible design direction is to consider how to adjust such existing interactions to better support active reading. However, our studies reveal a rich variety of actions, many of which are typically not or only partially supported in interactive visualizations. The behaviors identified during these studies suggest that there is room to extend our digital environments to come closer to supporting the diversity of activities that exist with paper. Implications for future research include enhancing visualizations with freeform annotation support, providing active reading for different types of visualizations, and exploring the possibility of supporting personalized active reading.

An important implication for the design of new visualization systems is considering and supporting naturally occurring active reading actions. Since teaching active reading of text has been shown to improve immediate comprehension and subsequent reading comprehension [55], it is possible to envision that active reading visualization environments might also come closer to the goal of amplifying a visualization reader's cognition by assisting and guiding active reading tasks.

8 LIMITATIONS AND FUTURE WORK

This work contributes to our understanding and characterization of visualization reading tasks. Through these studies, we identified low-level observable actions that we grouped into higher level goals. Because of the multiplicity of possible visualizations and the variation in tasks that engage visualization reading, we may only have discovered a subset of the possible visualization reading actions and goals. Creating a fuller reading task taxonomy is an important future research direction.

While our studies indicate that it is likely that the potential benefits of active reading may translate to visualization comprehension, they just scratch the surface of the investigations needed to fully understand the possible benefits. For example, both of our studies used just one type of visualization, node-link diagrams. Additionally, to keep the studies to reasonable lengths for our participants, a limited set of tasks were used. It is quite likely that different actions might be associated with different visualizations. The active reading techniques uncovered in S1 simply set the stage for the beginning of an active reading theory for visualization. Exploring the extent to which these ideas generalize to other visualization types and tasks requires considerable future work.

Supporting a wide variety of styles within a single visualization may be important because just one active reading technique is unlikely to fit all readers. However, just as there are sets of active text reading techniques taught to students, it may be possible to teach visualization readers sets of strategies, suggesting the possibility of a toolbox of techniques to choose from when reading. The spectrum of low level actions, together with the list of active reading goals, provides a framework that the Vis community can use to uncover more active reading techniques.

9 CONCLUSIONS

In this research, we proposed the concept of active reading of visualizations as a parallel to active reading of text, defining active reading of visualizations as *the purposeful, engaged reading of a visualization that combines internal and external reading strategies with available interactions to gain deeper comprehension*. Exploring the concept of active reading of visualizations is a step towards better supporting the needs of readers of visualizations. Our qualitative exploratory study results show that, in a paper-based visualization context, people perform a large number of physical actions with the visualizations to support their reading process. Our quantitative study results demonstrate that the support for active reading actions that can be achieved with a freeform annotation layer in an interactive environment (including altering views by making marks on or near visualizations) leads to measurable improvements in the accuracy of reading visualizations. Together our studies led to five high-level goals for active reading of visualizations: recognizing, tracking, reorganizing, decoding and analyzing.

Comprehensive exploration of active reading of visualizations has all indications of continuing to be fascinating. This work unveils many future directions including: exploring whether teaching how to read visualizations actively will bring benefits similar to those with text; how to effectively support positioning actions for active reading of visualizations; investigating visualization parallels for the syntopical reading level for text; and discovering what is appropriate support for active reading in digital environments. This initial exploration of active reading of visualizations suggests that it is worth investigating new ways to support readers of visualizations such as by offering access to a variety of low-level active reading actions.

10 SUPPLEMENTAL MATERIAL

In the interest of replicability, we provide additional materials about the studies on the accompanying website: <http://innovis.cpsc.ualgary.ca/supplemental/Active-Reading-of-Visualizations>

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their thoughtful comments and feedback. This research was supported in part by: Alberta Innovates - Technology Futures (AITF); Natural Sciences and Engineering Research Council of Canada (NSERC); and SMART Technologies.

REFERENCES

- [1] Active reading — Skills for OU study — Open University, August 2016.
- [2] A. Adler, A. Gujar, B. L. Harrison, K. O'Hara, and A. Sellen. A diary study of work-related reading: Design implications for digital reading devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '98, pp. 241–248. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1998. doi: 10.1145/274644.274679
- [3] M. J. Adler and C. Van Doren. *How to read a book: The classic guide to intelligent reading*. Simon and Schuster, 1972.
- [4] B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J.-D. Fekete. Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pp. 483–492. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2470724
- [5] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *Proceedings of the Proceedings of the 2005 IEEE Symposium on Information Visualization*, INFOVIS '05, pp. 15–. IEEE Computer Society, Washington, DC, USA, 2005. doi: 10.1109/INFOVIS.2005.24
- [6] American Psychological Association. *The Publication manual of the American psychological association*. Washington, DC, 6th ed., 2013.
- [7] C. Andrews, A. Endert, and C. North. Space to think: Large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pp. 55–64. ACM, New York, NY, USA, 2010. doi: 10.1145/1753326.1753336
- [8] J. Bertin. *Graphics and Graphic Information Processing*. De Gruyter, 1981.
- [9] J. Bertin. *Semiology of Graphics*. University of Wisconsin Press, Wisconsin, US, 1983.
- [10] M. Bostock, V. Ogievetsky, and J. Heer. D³: Data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, Dec 2011. doi: 10.1109/TVCG.2011.185
- [11] J. Boy, L. Eveillard, F. Detienne, and J.-D. Fekete. Suggested interactivity: Seeking perceived affordances for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):639–648, Jan 2016. doi: 10.1109/TVCG.2015.2467201
- [12] J. Boy, R. A. Rensink, E. Bertini, and J.-D. Fekete. A principled way of assessing visualization literacy. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1963–1972, Dec 2014. doi: 10.1109/TVCG.2014.2346984
- [13] N. Boyles. Closing in on close reading. In M. Scherer, ed., *On Developing Readers: Readings from Educational Leadership*, EL Essentials, chap. 9, pp. 89–99. ASCD, 2016.
- [14] P. Brandl, C. Richter, and M. Haller. Nicebook: Supporting natural note taking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pp. 599–608. ACM, New York, NY, USA, 2010. doi: 10.1145/1753326.1753417
- [15] W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554, 1984.
- [16] C. Collins, S. Carpendale, and G. Penn. Docuburst: Visualizing document content using language structure. *Computer Graphics Forum (Proc. of the Eurographics/IEEE-VGTC Symposium on Visualization (EuroVis))*, 28(3):1039–1046, 2009. doi: 10.1111/j.1467-8659.2009.01439.x
- [17] G. Cumming and S. Finch. Inference by eye: Confidence intervals and how to read pictures of data. *American Psychologist*, 60(2):170, 2005.
- [18] E. Dimara, A. Bezerianos, and P. Dragicevic. The attraction effect in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):471–480, Jan 2017. doi: 10.1109/TVCG.2016.2598594
- [19] P. Dragicevic. Fair statistical communication in hci. In J. Robertson and M. Kaptein, eds., *Modern Statistical Methods for HCI*, pp. 291–330. Springer International Publishing, Cham, 2016. doi: 10.1007/978-3-319-26633-6_13
- [20] C. Dunne and B. Shneiderman. Motif simplification: Improving network visualization readability with fan, connector, and clique glyphs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pp. 3247–3256. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2466444
- [21] M. Ghoniem, J.-D. Fekete, and P. Castagliola. A comparison of the readability of graphs using node-link and matrix-based representations. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*, pp. 17–24, 2004. doi: 10.1109/INFVIS.2004.1
- [22] P. Goffin, W. Willett, J.-D. Fekete, and P. Isenberg. Exploring the placement and design of word-scale visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2291–2300, Dec 2014. doi: 10.1109/TVCG.2014.2346435
- [23] L. Grammel, M. Tory, and M. A. Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, Nov 2010. doi: 10.1109/TVCG.2010.164
- [24] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pp. 1029–1038. ACM, New York, NY, USA, 2007. doi: 10.1145/1240624.1240781
- [25] N. Henry, A. Bezerianos, and J.-D. Fekete. Improving the readability of clustered social networks using node duplication. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1317–1324, Nov. 2008. doi: 10.1109/TVCG.2008.141
- [26] N. Henry and J.-D. Fekete. Matlink: Enhanced matrix visualization for analyzing social networks. In C. Baranauskas, P. Palanque, J. Abascal, and S. D. J. Barbosa, eds., *Human-Computer Interaction – INTERACT 2007 Proceedings, Part II*, pp. 288–302. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-74800-7_24
- [27] K. Hinckley, X. Bi, M. Pahud, and B. Buxton. Informal information gathering techniques for active reading. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pp. 1893–1896. ACM, New York, NY, USA, 2012. doi: 10.1145/2207676.2208327
- [28] U. Hinrichs, S. Forlini, and B. Moynihan. Speculative practices: Utilizing infovis to explore untapped literary collections. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):429–438, Jan 2016. doi: 10.1109/TVCG.2015.2467452
- [29] M. Hong, A. M. Piper, N. Weibel, S. Olberding, and J. Hollan. Micro-analysis of active reading behavior to inform design of interactive desktop workspaces. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces, ITS '12*, pp. 215–224. ACM, New York, NY, USA, 2012. doi: 10.1145/2396636.2396670
- [30] J. Hullman, N. Diakopoulos, E. Momeni, and E. Adar. Content, context, and critique: Commenting on a data visualization blog. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & #38; Social Computing, CSCW '15*, pp. 1170–1175. ACM, New York, NY, USA, 2015. doi: 10.1145/2675133.2675207
- [31] S. Huron, Y. Jansen, and S. Carpendale. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2102–2111, 2014.
- [32] P. Isenberg, A. Tang, and S. Carpendale. An exploratory study of visual information analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, 2008.
- [33] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges. In R. Borgo, F. Ganovelli, and I. Viola, eds., *Eurographics Conference on Visualization (EuroVis) - STARs*. The Eurographics Association, 2015. doi: 10.2312/eurovisstar.20151113
- [34] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pp. 2593–2602. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2481359
- [35] M. Kay and J. Heer. Beyond weber'law: A second look at ranking visualizations of correlation. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):469–478, Jan 2016. doi: 10.1109/TVCG.2015.2467671
- [36] K. N. Kirby and D. Gerlanc. Bootes: An r package for bootstrap confidence intervals on effect sizes. *Behavior Research Methods*, 45(4):905–927, 2013. doi: 10.3758/s13428-013-0330-5
- [37] D. Kirsh. Thinking with external representations. *AI & SOCIETY*, 25(4):441–454, 2010. doi: 10.1007/s00146-010-0272-8
- [38] B. C. Kwon and B. Lee. A comparative evaluation on online learning approaches using parallel coordinate visualization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 993–997. ACM, 2016.
- [39] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry. Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI Workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization*, BELIV '06, pp. 1–5. ACM, New York, NY, USA, 2006. doi: 10.1145/1168149.1168168
- [40] S. Lee, S.-H. Kim, Y.-H. Hung, H. Lam, Y.-a. Kang, and J. S. Yi. How do people make sense of unfamiliar visualizations?: A grounded model

- of novice's information visualization sensemaking. *IEEE transactions on visualization and computer graphics*, 22(1):499–508, 2016.
- [41] N. Mahyar and M. Tory. Supporting communication and coordination in collaborative sensemaking. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1633–1642, Dec 2014. doi: 10.1109/TVCG.2014.2346573
- [42] F. Matulic and M. C. Norrie. Supporting active reading on pen and touch-operated tablets. In *Proceedings of the International Working Conference on Advanced Visual Interfaces, AVI '12*, pp. 612–619. ACM, New York, NY, USA, 2012. doi: 10.1145/2254556.2254669
- [43] N. McCurdy, J. Lein, K. Coles, and M. D. Meyer. Poemage: Visualizing the sonic topology of a poem. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):439–448, Jan 2016. doi: 10.1109/TVCG.2015.2467811
- [44] J. E. McGrath. Human-computer interaction. chap. Methodology Matters: Doing Research in the Behavioral and Social Sciences, pp. 152–169. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995.
- [45] D. S. McNamara, I. B. Levinstein, and C. Boonthum. iSTART: Interactive strategy training for active reading and thinking. *Behavior Research Methods, Instruments, & Computers*, 36(2):222–233, 2004.
- [46] F. Moretti. *Graphs, Maps, Trees: Abstract Models for a Literary History*. Verso, 2005.
- [47] M. R. Morris, A. J. Brush, and B. R. Meyers. Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP '07. Second Annual IEEE International Workshop on*, pp. 79–86, Oct 2007. doi: 10.1109/TABLETOP.2007.12
- [48] C. Neustaedter and P. Sengers. Autobiographical design in HCI research: Designing and learning through use-it-yourself. In *Proceedings of the Designing Interactive Systems Conference, DIS '12*, pp. 514–523. ACM, New York, NY, USA, 2012. doi: 10.1145/2317956.2318034
- [49] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, May 2006. doi: 10.1109/MCG.2006.70
- [50] K. O'Hara. Towards a Typology of Reading Goals. Technical report, Rank Xerox Research Centre, Cambridge, U.K., 1996.
- [51] J. Pearson, G. Buchanan, and H. Thimbleby. *Designing for Digital Reading*. Morgan & Claypool, 2013. doi: 10.2200/S00539ED1V01Y201310ICR029
- [52] C. Perin, P. Dragicevic, and J. Fekete. Revisiting bertin matrices: New interactions for crafting tabular visualizations. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2082–2091, 2014. doi: 10.1109/TVCG.2014.2346279
- [53] C. Perin, F. Vernier, and J.-D. Fekete. Interactive horizon graphs: Improving the compact visualization of multiple time series. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, pp. 3217–3226. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2466441
- [54] M. Pohl, M. Schmitt, and S. Diehl. Comparing the readability of graph layouts using eyetracking and task-oriented analysis. In *Proceedings of the Fifth Eurographics Conference on Computational Aesthetics in Graphics, Visualization and Imaging*, Computational Aesthetics'09, pp. 49–56. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2009. doi: 10.2312/COMPAESTH/COMPAESTH09/049-056
- [55] C. Porter-O'Donnell. Beyond the yellow highlighter: Teaching annotation skills to improve reading comprehension. *English Journal*, pp. 82–89, 2004.
- [56] S. Posavec. Writing without words, 2006.
- [57] M. N. Price, B. N. Schilit, and G. Golovchinsky. Xlibris: The active reading machine. In *CHI 98 Conference Summary on Human Factors in Computing Systems, CHI '98*, pp. 22–23. ACM, New York, NY, USA, 1998. doi: 10.1145/286498.286510
- [58] B. Saket, P. Simonetto, S. G. Kobourov, and K. Börner. Node, node-link, and node-link-group diagrams: An evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2231–2240, Dec 2014. doi: 10.1109/TVCG.2014.2346422
- [59] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):443–456, July 2005. doi: 10.1109/TVCG.2005.53
- [60] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, Nov 2010. doi: 10.1109/TVCG.2010.179
- [61] I. Spence. Visual psychophysics of simple graphical elements. *Journal of Experimental Psychology: Human Perception and Performance*, 16(4):683, 1990.
- [62] A. L. Strauss and J. Corbin. *Basics of qualitative research: Grounded theory procedures and techniques*. Sage publications Newbury Park, CA, 1990.
- [63] C. S. Tashman and W. K. Edwards. Active reading and its discontents: The situations, problems and ideas of readers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, pp. 2927–2936. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979376
- [64] C. S. Tashman and W. K. Edwards. Liquidtext: A flexible, multitouch environment to support active reading. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, pp. 3285–3294. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979430
- [65] B. Tversky, J. B. Morrison, and M. Betrancourt. Animation: can it facilitate? *International Journal of Human-Computer Studies*, 57(4):247–262, 2002.
- [66] C. Wacharamanotham, K. Subramanian, S. T. Völkel, and J. Borchers. Statsplorer: Guiding novices in statistical analysis. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pp. 2693–2702. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702347
- [67] J. Walny, S. Huron, and S. Carpendale. An Exploratory Study of Data Sketching for Visual Representation. *Computer Graphics Forum*, 2015. doi: 10.1111/cgf.12635
- [68] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
- [69] W. Willett, S. Ginosar, A. Steinitz, B. Hartmann, and M. Agrawala. Identifying redundancy and exposing provenance in crowdsourced data analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2198–2206, Dec 2013. doi: 10.1109/TVCG.2013.164
- [70] W. Willett, P. Goffin, and P. Isenberg. Understanding digital note-taking practice for visualization. *IEEE Computer Graphics and Applications*, 35(4):38–51, July 2015. doi: 10.1109/MCG.2015.52
- [71] W. Willett, B. Jenny, T. Isenberg, and P. Dragicevic. Lightweight relief shearing for enhanced terrain perception on interactive maps. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pp. 3563–3572. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702172
- [72] K. Xu, C. Rooney, P. Passmore, D.-H. Ham, and P. H. Nguyen. A user study on curved edges in graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2449–2456, Dec. 2012. doi: 10.1109/TVCG.2012.189
- [73] D. Yoon, N. Chen, and F. Guimbretière. Texttearing: Opening white space for digital ink annotation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, UIST '13*, pp. 107–112. ACM, New York, NY, USA, 2013. doi: 10.1145/2501988.2502036
- [74] J. Zhao, M. Glueck, F. Chevalier, Y. Wu, and A. Khan. Egocentric analysis of dynamic networks with egolines. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, pp. 5003–5014. ACM, New York, NY, USA, 2016. doi: 10.1145/2858036.2858488