Interpretive Impacts of Text Visualization: Mitigating Political Framing Effects

ERIC P. S. BAUMER, Lehigh University
JAIME SNYDER, University of Washington
GERI K. GAY, Cornell University

Information visualizations are often evaluated as a tool in terms of their ability to support performance of a specific task. This article argues that value can be gained by instead evaluating visualizations from a communicative perspective. Specifically, it explores how text visualization can influence the impacts that framing has on the perception of political issues. Using data from a controlled laboratory study, the results presented here demonstrate that exposure to a text visualization can mitigate framing effects. Furthermore, it also shows a transfer effect, where participants who saw the visualization remained uninfluenced by framing in subsequent texts, even when the visualization was absent. These results carry implications for the methods used to evaluate information visualization systems, for understanding the cognitive and interpretive mechanisms by which framing effects occur, and for exploring the design space of interactive text visualization.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; **Visualization design and evaluation methods**; *Social engineering (social sciences)*;

Additional Key Words and Phrases: Text visualization, visualization design

ACM Reference format:

Eric P. S. Baumer, Jaime Snyder, and Geri K. Gay. 2018. Interpretive Impacts of Text Visualization: Mitigating Political Framing Effects. *ACM Trans. Comput.-Hum. Interact.* 25, 4, Article 20 (August 2018), 26 pages.
<https://doi.org/10.1145/3214353>

1 INTRODUCTION

“We change our tools and then our tools change us.” – Harold Innis [66, 122]

Technologies associated with the digital representation of text have evolved rapidly in recent decades, from hypertext markup to text-based data visualization. Currently, people with varying levels of technical sophistication are able to create novel re/representations of text with relative speed and ease. Customizations to the presentation of text can include graphic or stylistic changes (e.g., typeface, color, and kerning) as well as changes that fundamentally alter a text’s structure, spatial configuration, and interaction capabilities (e.g., hyperlinks, word clouds, interactive animations, and streaming data inputs). In the past, these customizations occurred primarily through

This work is supported by the National Science Foundation, under grant IIS-1110932.

Authors’ Addresses: E. P. S. Baumer, Computer Science & Engineering, 379 Packard Lab, 19 Memorial Blvd. W., Lehigh University, Bethlehem, PA 18015; email: ericpsb@lehigh.edu; J. Snyder, Information School, University of Washington, Mary Gates Hall, Suite 370, Box 352840, Seattle, WA 98195-2840; email: jas1208@uw.edu; G. K. Gay, Information Science Department, 219 Gates Hall, Cornell University, Ithaca, NY 14853; email: gkg1@cornell.edu.

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 ACM 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.

© 2018 ACM 1073-0516/2018/08-ART20 \$15.00

<https://doi.org/10.1145/3214353>

techniques, such as hypertext markup language (HTML) and cascading style sheets (CSS). More recently, novel tools for visualizing text, such as D3 [14] and Gephi [5], enable users to create and manipulate complex interactive text displays. In spite of the increasing ease with which we can manipulate the material and visual presentation of text using digital tools, few recent human computer interaction studies have been conducted to better understand potential interpretive impacts of interactive visual representations of text.

Particularly in terms of tools that handle text as data and enable users (especially researchers) to interactively visualize a corpus, designers, and developers often focus on how well an application helps an individual accomplish a particular analytic task [69, 76]. However, the impacts of these visualizations often transcend purely analytic contexts. Increasingly, visual representations of text (and/or other data) are used in communicative settings, such as the popular press [39, 83]. A different set of questions emerges regarding use of interactive text visualization in communicative contexts. How do visualizations influence interpretation and sense-making in comparison to prose? What is the nature of the relationship among the affordances of spatial display, the interactivity of different text elements, and the meanings conveyed by a text? What are the interpretive consequences and impacts of these relationships? Furthermore, to what extent do the interpretive effects of text visualizations transfer to situations where a visualization is absent?

These questions apply not only to graphic design audiences but also represent core issues in the design of interactive systems for the visual display of text. Currently, text visualizations are typically created using digital tools and often include filters, zooming, highlights, search, click-and-drag, and other means for viewers to manipulate the representation of data. Understanding the communicative and interpretive impacts of these affordances is essential to informing the design of interactive text visualizations.

To explore these questions, this article reports on a study comparing the interpretation of text in conventional prose form with that of two different visualizations. These include a standard word cloud visualization, chosen for its popularity and perceived accessibility [37, 83], and a custom *word web* visualization, which combines some of the co-occurrence structure of prosodic text, the non-linearity of word clouds, and an interactive browsing component reminiscent of hypertext.

Our study assesses how these different representations of text impact the communicative phenomenon of framing. In sociology and political communication, framing involves the selection, emphasis, and presentation of information, often through subtle linguistic differences, which shape our understanding of a situation [23, 27, 32, 34, 35, 45]. Significant prior work has shown that different framings of a given issue can significantly impact subsequent self-reported political opinions on that issue [2, 24, 25, 42, 65, 67, 71, 74, 77].

The results presented here show that visual representations of text significantly influence interpretation in communicative settings. In addition to the visual and interactive format of the presentation impacting the interpretation of the text upon initial exposure, we also saw a transfer effect. Participants who saw the visualizations demonstrated similar interpretive patterns in subsequent situations, even though visualizations were no longer present and they were being asked to interpret a completely different text.

These results provide two sets of implications: implications for future research and design, as well as more theoretical considerations. First, text in prose form may more effectively make use of subconscious rhetorical influences, such as that from framing, by not highlighting them. Conversely, visualizations of text may provide a form of inoculation (albeit temporary) to such subconscious influences. Designers of visual text representations may benefit from considering these varying impacts on interpretation.

Second, both visualizations tested here used a spatial representation, one of which was based on co-occurrence. Both similarly impacted interpretive processes, although only the effect of the

layout based on co-occurrence was statistically significant. This finding suggests that how a visualization structures possible interactions with a text (linear, hierarchical, spatial, and so on.) may play key roles in the interpretation of text visualization. Furthermore, the observed transfer effects indicate that the interpretive processes afforded by visualizations can to some extent become internalized, i.e., they may be learnable. Put differently, these results show how “our tools change us.”

2 RELATED WORK

2.1 The Digital Display of Text

Disrupting normative conventions can provoke people to see things differently. To understand these differences, it is necessary to consider the ways that material expression signals interpretive convention [48, 60]. According to social semiotic theories of modality, the form or format of an expression has a key role in the communicative function of the resulting message [49, 50]. For example, choices made regarding the display of text can result in both aesthetic cues (“fancy” lettering on a wedding invitation) and normative prompts regarding functional expectations (a magazine spread versus a dictionary page). Further, empirical studies have investigated how typography can also impact reading time, memory, and even mood [15, 54, 70]. In many ways, this study provides empirical support for insights accumulated by typographers and information designers over years of practical experience.

Furthermore, this study builds on those insights by extending them to interactive text. In 1963, Ted Nelson (first published in [59]) marked a new era in the history of the written word by introducing the terms *hypermedia* and *hypertext* to refer to the non-sequential flow of ideas made possible by computers and digital displays [58]. According to Nelson, the interactions afforded by hyperlinking would be more akin to how the human mind naturally structures ideas than standard written text; hypertext would free us from the constraints of ordinary sequential writing. The original dream of a decentralized, sprawling hypertext has morphed into the more stabilized reality of today’s blogs, social networks, and crowd-edited encyclopedias [47]. However the plastic, associative, and non-linear structure of hypermedia has had a lasting impact on the ways that we interact with text on the web.

Since the advent of hypermedia, digital environments have continued to provide novel opportunities for experimenting with how information is represented. However, as early experiments in hypertext storytelling show us, some disruptions are more generative than others. For example, typographers in the 1990s experimented with layouts that radically compressed the space between lines of text and introduced syncopating groupings of letters and words violating the traditional grid underlying most designed pages [81]. The results were often edgy and sometimes interesting, but rarely readable.

Today, computational tools provide designers and users the ability to manipulate both the style and structure of text with increasing ease. For example, the latest wave of web-based visualization tools, such as D3 [14] or Gephi [5], has introduced new possibilities for flexible and extensible presentation of text and other forms of data. Our study shows that experimenting with the representation of text has not only *aesthetic* but also *interpretative* implications. We chose to manipulate spatial and proxemic features to highlight linguistic aspects of political framing. Other visual encodings such as variations in length, size, or orientation could be applied to a corpus in order to explore patterns of language related to sentiment analysis [62, 90], racial bias [10, 80], or gender identity [4].

The point we highlight here is that choices made regarding the representation of text can impact mechanisms of interpretation (such as, in our case, mitigating the effects of framing). This

highlights the need for users and designers of text visualization tools to be aware of the ways that decisions made regarding the structural and stylistic representation of texts, whether for analysis or communication, can impact interpretive outcomes.

2.2 Text Visualization

Text visualizations are often designed to highlight patterns and trends in language that would otherwise be difficult to identify, thereby reducing cognitive load and enhancing analytic efficiency [87]. Creation of text visualizations typically involves applying techniques of text mining and natural language processing to represent a corpus as a set of quantitative characteristics, such as word frequency or patterns of co-occurrence [18]. Recently, new approaches to text visualization have emerged that reflect more nuanced aspects of language use. For example, discourse parsers help analysts extend beyond frequency and co-occurrence to identify trends associated with more complex language structures [93]. Authorship analysis of online content has become an important dimension of security-related intelligence work; a recently developed prototype system automatically detects distinctive linguistic features of authors at the message level in order to support the work of these analysts [8]. DAVIEWER [93], an interactive visualization system, helps computational linguists explore, compare, evaluate, and annotate the results of multiple discourse parsers.

Topic detection has also been recognized as a challenging aspect of automated text processing [12, 13], as the main focus of a corpus can shift and evolve over time and across documents [1]. TIARA is an interactive text visualization system that is designed to help users explore and analyze the evolution of topics contained in large bodies of text using time-sensitive keywords [86]. This system is intended to help analysts better understand how the text has been summarized through the keyword detection process. TextFlow is another tool for visualizing the evolution of topics in large text collections, such as newspaper reports or blog posts, by building on dynamic topic mining techniques [19]. The visual component of the TextFlow tool is designed to convey complex relationships among topic evolution trends, critical events, and keyword correlations.

In these and similar cases, the complexity of the text visualization corresponds closely with the degree of linguistic expertise possessed by the target users [64]. More sophisticated users benefit from more complex visual representations of nuanced language structures and meanings. It follows that communication-driven use cases benefit from visual encoding that affords interpretative accessibility and comprehensibility for a broad audience. This was an important consideration in the design of the text visualizations for our study. The participants in this study were not experts in discourse linguistics. They were not expected to be seeking or even aware of framing mechanisms in the text prior to having them highlighted in the two visualization conditions. Thus, in order to assess the interpretive effects of the visualization conditions in our study, we needed to ensure that the form and format of text visualizations were accessible and comprehensible to our target users (discussed further below).

Work emerging from the information visualization community highlights the importance of evaluating the effectiveness of visualizations from multiple perspectives, not just the performance of a discrete task. In recent years, there has been growing recognition in this community that traditional approaches to evaluation focused on time and error metrics may not always adequately measure the range of cognitive and communicative activities associated with working with visual representations of information [38, 43, 61]. For example, Isenberg et al. [44] have advocated for a grounded approach to evaluation that accounts for the situated nature of visualization practices and introduces qualitative methods for understanding contextual aspects of use early in the visualization design cycle. Techniques have also been proposed and developed for measuring insight generation [61, 91], sense-making processes [26], and narrative interpretation [42]. Dörk et al. discuss the ways that visualizations have become enmeshed within cultural communication,

arguing that this integration makes it essential for designers to systematically consider the values and intentions of both themselves and the consumers of their work [22].

Our study contributes to this work by re-contextualizing the potential benefits of a popular form of text visualization, the word cloud, as a means for supporting awareness of conceptual framing [27] in political texts. We employ evaluation methods that shift focus away from time and errors associated with the completion of specific interpretive tasks. Rather than primarily viewing word clouds as analytic tools, a function that has already been identified as a problematic for a variety of reasons [83], we instead capture evidence of how exposure to text visualizations can impact the interpretive influence of framing effects [17]. Framing, discussed next, represents an important part of the interpretive process.

2.3 Framing and Frame Reflection

Scholars from multiple disciplines have brought different perspectives to bear on understanding how framing can influence public opinion [17, 27, 35, 82]. Generally speaking, a frame consists of a variety of linguistic and paralinguistic components, including “keywords, stock phrases, stereotype images, sources of information” [27:52], “metaphors, exemplars, catchphrases, depictions, visual images,” [32:3] and other linguistic and conceptual devices. Discursive frames give meaning to observations of the world; “[facts] take on their meaning by being embedded in a frame [...] that organizes them and gives them coherence” [31:157]. Collectively, a frame provides an interpretive lens or “package” [32] through which to perceive and make sense of facts and events. This article draws primarily on psychological and cognitive accounts of framing from communication and political science [24, 25, 65, 74]. Doing so allows us to operationalize framing in a way that can be useful in terms of design and evaluation, as will be shown below.

Political communication research demonstrates how framing can influence public opinion on a variety of topics. For instance, self-identified conservatives are less likely to say they believe in “global warming” than in “climate change,” even though the two are often used synonymously [74]. Referring to same-sex marriage as “homosexual marriage” garners more polarized responses than the label “gay civil unions” [65]. Support for a national healthcare system is greater when healthcare is framed in terms of equal access to medical treatment than when framed in terms of costs associated with access [24]. One concern regarding unmediated framing effects, then, is the potential for political elites to exercise undue influence over public opinion by strategically choosing the frames that shape political discourse [2, 33].

Effortful, deliberate cognition, such as providing justifications for decisions or conversing with groups holding heterogeneous opinions, can mitigate framing effects [25, 40, 41, 57, 89]. Related research has shown that it may be beneficial to draw attention to, and encourage critical thinking and reflection about, the frame(s) offered for a given issue. Schön and Rein [73] argue for the value of *frame reflection*, wherein “assumptions, views of the world, and values that have heretofore remained in the background, giving shape to foreground inquiry but keeping, as it were, to the shadows, become foreground issues, open to discussion and inquiry in their own right” [68:94]. One means of enabling frame reflection involves simply using typographic methods such as bolding or underlining to draw attention or to highlight the language invoking a frame [7]. However, a challenge with such an approach is determining exactly which words or phrases to highlight. Despite a robust body of scholarship in the area, “straightforward guidelines on how to identify [...] a frame in communication do not exist” [17:106].

Visualizations offer a wider range of techniques for encoding relational characteristics – linearly, hierarchically, spatially, and so on. Because of this, they are often deployed in situations requiring individuals to complete comparative tasks, typically referred to as “decision-making.” In recent years, researchers have begun to explore the influence of visualizations on mechanisms

of cognitive bias that can effect judgment formation during decision-making processes [20, 92]. The types of choices required of participants in these studies are distinct from the interpretive and communicative practices described above. However, this work does bolster the notion that addressing the rich sociotechnical contexts within which digitally mediated text displays exist is essential for designing new methods for evaluating their impacts.

Our approach to representing text may similarly mitigate framing effects in communicative contexts with the added benefit of not relying on assumptions about specific possible frame-invoking terms. We argued above that text visualization can shape and alter the processes of interpretation that are so important to effective communication. We have also shown that framing effects are bound up with these processes of interpretation. Furthermore, prior work provides evidence that frames usually have their impact subconsciously [32], and that more effortful, deliberate cognition (not necessarily explicit awareness of specific linguistic mechanisms) is associated with mitigation of framing effects [25, 40, 41, 57, 89]. Therefore, the study presented here extends previous work on frame mitigation [7] by testing the hypothesis that the text visualization's ability to surface subtle linguistic patterns that pervade a text may thus draw conscious attention to those patterns in a way that impacts the interpretive processes by which frames have their effect.

Thus, a two-fold question arises. First, does the use of text visualization in a communicative context (as opposed to an analytic one) impact interpretive processes in a way that mitigates framing effects? Second, do mitigating effects occur only in the presence of the visualization, or might the alternative interpretive processes transfer to other situations where the visualization is absent?

3 STUDY MATERIALS AND VISUALIZATION DESIGN

In order to focus on how communicative uses of visualization impact interpretive processes, we wanted to select for study visualization styles that were common in such communicative settings. We determined that a word cloud would offer a suitable text visualization approach that could represent verbal content in text, while also drawing on a robust and familiar set of visual conventions to disrupt the typical linear representation of prose. Despite their “glaring theoretical problems” [83:58], use of tag and word clouds has proliferated among vernacular audiences [78]. For our purposes, the accessibility and simplicity of presentation provided by the word cloud format was highly desirable in order to engage participants. If text visualizations surface linguistic patterns in an accessible manner, there should be a reasonable expectation, based on the political communication literature, that these visual displays would likely influence the interpretive processes associated with framing effects by supporting more reflective engagement with the text.

As described above, framing is instantiated by various patterns of language stemming from associations among words, not simply frequency. Catchphrases [32], metaphors [32], stereotypes [27], and other similar invokers of framing all occur through consistent reuse of a fixed set of terms in close proximity. Identifying and drawing attention to such patterns of language requires at least two important considerations.

First, such patterns often involve inter-dependencies among different types of words, including both seemingly innocuous stopwords [55] – high frequency determiners, articles, pronouns, and so on – and more “loaded” or “poetic” terms that are much more open to interpretation [79]. For example, consider the phrases “atoms for peace” [32], “lipstick on a pig,” “affirmative action” [65], or “the war on terror.” The significance of these collections of words may not become readily apparent in visualizations such as conventional word clouds [83], which do not represent stopwords, co-occurrence, or textual structure.

Second, while there are schemas for identifying linguistic mechanisms related to sentiment [62], there does not exist a single formulaic template, or even a group of templates, that can be used consistently and reliably for identifying the language of framing [6, 16, 17]. As a result, traditional

visualizations of colocation or co-occurrence, such as phrase nets [36] or word trees [85], may not fully highlight relevant patterns or the connotative meanings thereof [79].

Put concisely, while conventional word clouds replace the linear structure of text with a pseudo-random spatial arrangement, doing so makes only the relative *frequencies* of words available for interpretation. Due to the importance of *structural arrangement* to invoking frames [28], we sought to create a visualization that would take advantage of the familiarity of word clouds for representing frequencies, but would also visually represent some elements of prosodic structure in a spatial fashion.

To that end, we created both customized static word clouds (see Figure 3 below) and dynamic word webs (see Figures 1 and 2 below) to highlight specific language choices that reflect framing effects. The study design treats exposure to the text visualization as an *experimental condition* rather than as an *analytic tool*. This enabled us to test (1) whether the visualizations provided increased awareness or mitigation of framing effects and (2) whether the effects provided by the visualizations transferred to other texts where there was no visualization support.

3.1 Word Webs Implementation

Word clouds create a pseudo-random spatial layout of the words that occur frequently in a text. Word size is a function of word frequency, often the natural logarithm of frequency. We generated word clouds using a freely available online tool [<https://www.jasondavies.com/wordcloud/>].

Word webs essentially add structure to word clouds. Similar to word clouds, word webs scale the size of each word in proportion to its frequency. However, word webs also use spatial layout to convey meaning. Words are arranged using a force-directed layout, where the strength of a connection between two words is based on both the *frequency* and the *proximity* of their co-occurrence. This representation provides information about both word counts and structural relationships within a text.

Word webs are generated by processing raw text using NLTK's [11] word- and punctuation-based tokenizer. The frequency of each word is calculated, excluding numbers, stopwords, and words less than three letters long.

We then calculate an association strength for all pairs of words w_1 and w_2 that occur in the same sentence at least once in the text. Psychometric judgments of pairwise word associations are best predicted using the harmonic mean of the number of words that occur between the pair, especially when using sparser data [84]. Thus, for words w_1 and w_2 , the association strength is

$$\text{AssocStr}(w_1, w_2) = \sum_{i=0}^p \frac{1}{1+i} * \text{count}(w_1, w_2, i)$$

where i is the number of words that occur between w_1 and w_2 and $\text{count}(w_1, w_2, i)$ is the number of times w_1 and w_2 appear with i words between them. Thus, every time two words appear immediately next to each other, we add 1 to their association strength. Every time the same two words appear with one intervening word we add 0.5, with two intervening words we add 0.3 and so on, i.e., a harmonic progression. Here, manual tuning led to an upper limit of $p = 4$ intervening words. Ultimately, word pairs appearing closer together with higher frequency will have greater association strength than those occurring together less often or with more words between them.

These association strengths are used as input to a force-directed graph of words, rendered using D3 [14]. Words with a stronger association strength (i.e., that frequently occur in close proximity) are rendered more closely together in the resulting graph. The weight (i.e., thickness) of the line connecting two terms also scales directly with the association strength. The lines, then, serve

two important functions. First, they combine with distance between words to show association strength via another channel. Second, since a force-directed layout can result in any two unconnected nodes appearing in close proximity, the lines indicate which distances among pairs of nodes are meaningful. Thus, the lines connecting words both help ease legibility issues with the randomness of a force-directed layout, and they allow users to explore neighborhoods of connections, as described below.

When first viewed, the visualization shows all words and lines in the web using a semi-transparent gray color. However, in dense graphs, it can be difficult to see individual connections [30]. Thus, we also include a neighborhood highlighting feature. When a user hovers her or his mouse cursor over a given word, that word along with the words connected to it turn red, and all other words fade slightly. The original text is displayed to the right of the word web (see Figure 1), but it is initially too transparent to read. When a user hovers over a word in the word web, all the sentences in which that word appear in the full text of the article are also highlighted in red.

This neighborhood highlighting enables iterative perusal of associations. For example, the viewer may notice that the word “disparities” occurs frequently in one text about healthcare (Figure 1, bottom). To determine what kind of disparities, s/he can hover the cursor to see associations with, among other terms, “racial” (Figure 2, top). Specifically, when s/he hovers over racial, s/he sees connections with “whites,” “minorities,” and “gap” (Figure 2, middle), indicating a potential gap between healthcare coverage for whites and minorities. To determine which minorities, the user again follows the associations to “poor,” “black,” and “men” (Figure 2, bottom).

Following this trail of associations, enabled by the word web interface, helps to reveal the article’s framing. That is, it highlights the specific kinds of disparities being acknowledged in the text and how the (in)equality is used to frame the discussion of healthcare.

Furthermore, this visualization reorganizes the text. It replaces the conventional linear structure with a network structure based on co-occurrence patterns. Both the word web graph and the text itself become navigable in a spatial, network-based fashion. Thus, word webs provide a point of comparison both with unstructured word clouds and linearly structured prose.

4 EVALUATION

To understand how text visualization might impact the interpretive processes involved in framing effects, we conducted a controlled laboratory study. The study design compared the effect of reading a text in prose form to reading the text using one of two visualizations. The results show that both visualizations mitigated observed framing effects.

The study design tested for subsequent transfer effects to new text where the visualization is no longer present. The results show evidence for a transfer effect. Participants who saw a visualization were less susceptible to framing effects in subsequent questions, even in the absence of the text visualization.

4.1 Methods and Materials

Participants were asked to read a short news article discussing a political topic. Participants were then asked to answer a series of questions related to (1) their own opinion about the topic, (2) the author’s opinion about the topic and (3) a series of unrelated texts presented after the initial exposure.

The study used a 2×3 between-subjects experimental design, with each participant randomly assigned to one condition for each of the two experimental manipulations. The first experimental manipulation varied the frame to which participants were exposed. The news article that participants read framed healthcare either in terms of cost (i.e., the economic burden of providing

Fig. 1. Screenshots of the word webs for two articles about healthcare, one using a cost framing (top) and one an equality framing (bottom). In the top, the user has hovered over “insurance.” In the bottom screen shot, the user has clicked on “disparities.”

healthcare for everyone) or in terms of equality (i.e., demographic disparities in terms of who has access to healthcare). These materials came from a previous study [24], which found that exposure to the equality framing was associated with increased support for national healthcare, and the cost framing was associated with decreased support. This aspect of the study design also follows a common approach in the framing literature where study participants are exposed to exactly one of a limited number of frames [2, 65, 67, 71, 74, 77, 82].

Fig. 2. Series of screenshots showing a user following a chain of associations in the *equality* framed article: from “disparities” to “racial” (top) to “minorities” (middle) to “poor” to “black” to “men” (bottom). These associations reveal how the article’s framing emphasizes racial minorities and economic disparities.

The second experimental manipulation varied the manner in which the article text was presented. In the control condition (*text only*), participants saw only the text of the article as prose. In the first experimental visualization condition (*word cloud*), participants saw a word cloud, generated using www.jasondavies.com/wordcloud (see Figure 3), with the text of the article displayed below the visualization. For the word cloud, the randomly assigned orientation angles

Table 1. Study Design. Six Experimental Conditions from Three Visualization Manipulations (Listed Across the Top) and Two Frames (Listed Down the Left Side)

Framing	Control (Text Only)	Visualization 1 (Word Cloud)	Visualization 2 (Word Web)
Cost	Cost Text	Cost Word Cloud	Cost Word Web
Equality	Equality Text	Equality Word Cloud	Equality Word Web

and saturation/luminance values help make each word visually distinct from neighboring words of a similar size. Participants were also provided with the following description of the word cloud:

To help you understand the article better, the following visualization shows the most frequent words in the article. The size of each word represents how often it occurs in the article. Please take a few moments to review this visualization.

In the second experimental visualization condition (*word web*), participants saw a word web visualization, as described above, along with the full text of the article. Table 1 summarizes the experimental design. Participants were also provided with the following description of the word web:

To help you understand the article better, the following visualization shows the most frequent words in the article. The size of each word represents how often it occurs in the article. Hovering your mouse cursor over a word will also show you the other words near which it appears. Words with a thicker line between them appear together more often and closer together, while words with thinner lines do not occur together as often or have other words that appear between them. Please take a few moments to review this visualization.

This description omits an interpretation of the distance between pairs of words for two reasons. First, the exact placement within the visualization is a function of D3's force directed layout. Thus, exact distance within the spatial layout may not perfectly align with the underlying co-occurrence data. This is why co-occurrence strength is encoded as the thickness of the line between pairs of words. Second, proximity is one of the few pre-attentive visio-cognitive relationships that humans unconsciously default to interpreting literally [30]. Thus, without prompting, subjects should interpret the distance between words as an indicator of relatedness.

In both visualization conditions, the graphical layout placed the visualization first, followed by the text. An earlier experiment with visualizations placed after the text found minimal differences across the visualization conditions, so the experimental materials were designed for primacy of the visualization. The word cloud was placed above the text, so the viewer saw the cloud first and then scrolled down to the text. The word web was placed to the left of the text, and the text was fully legible only after the viewer interacted with the visualization to highlight a word (see Figure 1). The word web required both the text and visualization to be visible simultaneously, hence the side-by-side lay out. The word cloud required more horizontal screen real estate to be legible, hence the above-below layout. In both cases, though, the visualization acted as the participant's primary exposure to the text.

Participants were required to spend at least 60 seconds reading the article or viewing the visualization. Participants were then asked the following question, taken from [24:435]:

Some people feel there should be a universal government insurance plan that would cover medical and hospital expenses for all citizens. Others feel that medical

and hospital expenses should be paid by individuals and through private insurance plans. Where would you place yourself on this scale?

Participants responded on a seven-point Likert scale, where one end of the scale indicated that “All expenses should be covered by private insurance plans,” while the other indicated that “All expenses should be covered by universal government insurance.” This question tested for framing effects, helping to determine whether framing effects observed in the control condition were mitigated with the visualizations.

Participants then saw a series of questions, some taken from other studies of framing effects and some of our own design. Each question had two variants, where each variant invoked one frame or another. Each participant saw only one variant of each question, with the variant selected randomly for each question independently of the variants seen for the other questions and of the participant’s experimental condition. Questions covered the following issues:

- Climate: Belief in *global warming* vs. belief in *climate change* [74].
- Free Speech: Whether the United States should *allow* speeches against democracy vs. whether the United States should *forbid* speeches against democracy [71].
- Gun Legislation: Support for *gun control* vs. support for *gun rights*.
- Same-Sex Marriage: Support for *homosexual marriage* vs. support for *gay civil unions* [65].
- Abortion: An *unborn child’s right* to life vs. a *mother’s right* to end her pregnancy.
- Economic Policy: Preferences between two programs when framed as impacting *unemployment* vs. impacting *employment* [67].

These questions were used to determine whether any impacts of the visualization on sensitivity to framing effects transferred to questions where the participants were not directly exposed to a visualization.

Participants were also asked to respond to an open-ended question about how they used the visualization (or, in the control condition, the text) when answering the questions. Finally, participants completed a series of demographic questions, including age, gender, frequency of political discussion, education level, vision (normal, corrected to normal, visually impaired), and political ideology.

4.2 Participants

We recruited students from a social science subject pool at a major US research university. Participants completed the study instrument online. We excluded incomplete responses. The distribution of completion time was highly skewed, with the large majority completing in less than 20 minutes but a few participants taking as long as several days. Since this was an online study, we suspect that those participants were multi-tasking, completing the study, and doing other activities simultaneously. Following Malhotra [56], we exclude those participants whose completion times were greater than one standard deviation above the mean, which in our case was 16.6 minutes.

This left, a total $N = 174$ for the main analysis. Browser meta-information indicated that participant display resolutions ranged from 1093×614 pixels to 1920×1200 pixels. Participants were mostly female (132, 75.9%), with two respondents indicating a non-binary gender. The sample skewed liberal (97, 55.7%), with only 35 participants (20.1%) indicating Slightly Conservative or Conservative; no participant reported an ideology of Very Conservative. Age ranged from 18 to 46 ($M = 20.0$, $Mdn = 20$). Most participants had either a high school degree (30, 17.2%), some college but no degree (104, 59.8%) or a bachelor’s degree (33, 19.0%). Table 2 shows a full cross-tabulation of demographics across the framing that participants saw and the experimental condition to which

Table 2. Participant Demographics Cross Tabulated by Framing and by Experimental Condition

			Framing		Experimental Condition		
			Cost	Equality	Text	Cloud	Web
		Totals	90	84	56	59	59
Gender	Female	132	65	67	49	39	44
	Male	40	24	16	7	20	13
	Non-binary	2	1	1	0	0	2
Ideology	Conservative	35	12	23	13	9	13
	Centrist	42	26	16	14	17	11
	Liberal	97	52	45	29	33	35
Education	High school graduate	30	14	16	11	9	10
	Some College	104	52	52	34	36	34
	Associates Degree	5	3	2	4	0	1
	Bachelor's Degree	33	19	14	7	14	12
	Graduate Degree	1	1	0	0	0	1
Age – mean (st. dev.)		20.0 (2.3)	20.2 (3.0)	19.9 (1.2)	19.9 (1.3)	20.0 (3.6)	20.1 (1.2)

they were assigned. Because all participants were at least high school graduates, Table 2 omits rows for education levels of some high school and less than 8th grade.

4.3 Results

4.3.1 Framing Effect Results. To examine how the visualizations influenced the cognitive and interpretive processes of framing, we first looked at responses to the question asking about the participant's support for national healthcare. To reiterate, our outcome variable was the participant's self-reported support for a national healthcare system on a seven-point ordinal Likert scale from "All expenses should be covered by universal government insurance" to "All expenses should be covered by private insurance plans" [24]. An ordinal logistic regression model was used to examine how the framing condition, the visualization condition, and self-reported ideology impacted support for national healthcare. We expected that the responses of participants in the control condition (*text only*) would provide evidence of a framing effect, but not those of participants in the two experimental visualization conditions (*word cloud* and *word web*).

Table 3 shows full details of the model, including odds ratios and *p*-values for each predictor. Odds ratios greater than one indicate increased likelihood of supporting national healthcare, and odds ratios less than one indicate decreased likelihood. For example, for each step toward the conservative end of the self-reported political ideology measure, participants are 0.477 times as likely (i.e., 2.096 times *less* likely) to support national healthcare. This shows a main effect for

Table 3. Ordinal Logistic Regression Predicting Participant's Level of Support for National Healthcare, Including Odds Ratios and p -Values

Predictor	Odds Ratio	p -Value
Cost framing	0.047	0.035*
Word Cloud	0.282	0.405
Word Web	0.285	0.283
Ideology	0.477	0.001**
Cost $\times$ Word Cloud	34.77	0.095
Cost $\times$ Word Web	35.67	0.054
Cost $\times$ Ideology	2.200	0.047*
Word Cloud $\times$ Ideology	1.183	0.669
Word Web $\times$ Ideology	1.558	0.159
Cost $\times$ Word Cloud $\times$ Ideology	0.372	0.099
Cost $\times$ Word Web $\times$ Ideology	0.354	0.043*

Interactions are denoted by an "x" between predictor names. For significance, * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

ideology, that conservatives tend to oppose national healthcare and liberals tend to support it ($p = 0.001$). Note that the reference level for experimental condition was set as Text Only (since it was the control condition). Thus, odds ratios for the Word Cloud and Word Web compare each with the Text Only condition. Similarly, the reference level for framing was set as Equality, thus odds ratios show the impact of the Cost framing as compared to the Equality framing. While these choices of reference level impact interpretation of individual coefficients and odds ratios, they do not impact the significance of individual predictors or the overall soundness of the model [63].

The results also show a significant main effect for the framing. Participants who saw the cost frame were less likely to support national healthcare ($p = 0.035$). This result replicates results from a prior study that used the same materials [24] and serves as an important manipulation check. An alternate model controlling for education, gender, frequency of political discussion, and exposure to diverse political opinions was also tested. This model included similar odds ratios and significance values for the Cost framing (OR = 0.057, $p = 0.053$), Ideology (OR = 0.478, $p = 0.002$), the interaction of Cost $\times$ Ideology (OR = 2.098, $p = 0.070$), and the three-way interaction Cost $\times$ Word Web $\times$ Ideology (OR = 0.349, $p = 0.047$). None of the additional control variables (education, gender, frequency of political discussion, and exposure to diverse political opinions) have a significant effect. Furthermore, the alternate model has a higher over AIC (581.91 vs. 574.63), suggesting these additional control variables do not add significantly more explanatory power. Thus, the effects documented in Table 3 persist when controlling for these additional variables, and the model in Table 3 provides a better overall fit.

A few significant interactions also occurred. To ease interpretation of these interactions, we consider a slightly simplified version of the analysis in which we examined the likelihood of a participant's opinion about national healthcare being opposed (Likert < 4), supportive (Likert > 4), or neutral (Likert = 4). Figure 4 shows how a participant's ideology predicts the probability of her or his position on national healthcare in each of the six combinations of framing condition (cost vs. equality) and experimental condition (text only vs. word cloud visualization vs. word web visualization).

This figure illuminates complex and important interactions among framing, experimental condition, and ideology. In the control (i.e., *text only*) condition, an interaction occurs between framing and ideology. With the equality framing, liberals are much more supportive of national healthcare

Fig. 4. Interactions among framing, visualization condition, and political ideology in support for national healthcare. A framing effect occurs in the Text Only condition; in the left column, slopes of the Oppose and Support lines are opposite for the two framings. However, that framing effect does not occur in either of the visualization conditions; in the middle and right columns, relationship between ideology and support for national healthcare does not differ between the two frame conditions.

and conservatives much more opposed (Figure 4, left column, bottom cell). However, with the cost framing, the relationship is the opposite: conservatives are more likely to support national healthcare and liberals more likely to oppose it (Figure 4, left column, top cell). As expected, this result confirms the presence of a framing effect in the control condition.

In the experimental conditions (i.e., *word cloud* and *word web*), the interaction between ideology and framing does not occur. Ideology has the same relationship with opinion about national healthcare regardless of the framing the participant sees. That is, the presence of the visualization eliminates the framing effect (Figure 4, center and right columns). Indeed, Table 3 shows a

significant three-way interaction among framing, ideology, and both the word web visualization ($p = 0.04$) and a near-significant interaction with the word cloud visualization ($p < 0.10$). These results do not arise from time on task, which does not differ significantly across visualization condition, the framing that participants saw, self-reported ideology, or education level. Furthermore, the only significant predictor for participants' perception of the author's opinion about national healthcare was the framing. That is, most participants (133, 76.4%) accurately understood the author's position regardless of the visualization condition, self-reported ideology, frequency of political discussion, or other control variables.

Participants' responses to the open-ended question about how they used the visualization (or text) shed some light on why these differences may have occurred. Those in the control condition often reported little influence from the text, for example, "I did not really use the text when answering the questions asked. I pretty much had pre-established opinions on the subjects" (P#148, cost framing). Participants' disavowal of the article's influence aligns with the subconscious influence of framing [32, 35, 52]. In the word cloud condition, many responses refer to treating the cloud as a summary: "The visualization was helpful in summing up the main points of the article and giving me a 'synopsis' of what it was going to be about" (P#22, cost framing). These responses suggest an approach to reading the article that focuses on summary points. Responses in the word web condition reflect a third approach: "I first clicked on the largest words and the corresponding connected words in order to get a gist of what the article was about. If I wanted to know about a certain topic I clicked on a word and read its corresponding paragraphs" (P#84, equality framing). These statements about consciously directed, non-linear approaches provide further evidence about how the visualizations impacted participants' process of interpreting the text.

4.3.2 Transfer Effect Results. The above results show that text visualizations can mitigate framing effects that would normally occur when reading the text in prose format. This section considers whether those effects transfer to a novel context. That is, after being exposed to a text visualization, do participants become less susceptible to framing effects in subsequent questions on different issues where a visualization has not been provided?

To test this question, after completing the first part of the protocol, participants in all conditions answered several opinion questions *in the absence of a visualization*. As described above, these questions were either drawn from prior literature on framing [65, 71, 74] or of our own design. To assess transfer effects, we performed similar analyses as above, determining whether participants in the word cloud and word web visualization conditions were more or less susceptible to the framing effects in these subsequent questions. In the interest of concision, we only report here responses to one question, along with a summary of results from the others.

The question analyzed here, one of our own design, framed gun legislation in terms of either "gun control" or "gun rights": "One prominent issue is gun <control/rights>. What is your personal opinion about stronger <controls on gun ownership/protections for the right of Americans to own guns>?" Responses were coded for their degree of support for gun ownership. We expected that participants would be more supportive of gun ownership when framed in terms of "gun rights" than in terms of "gun control," and that this framing effect would interact with self-reported ideology.

Indeed, in the control (*text only*) condition, a gun rights framing increased the likelihood of support for gun ownership among conservatives and decreased likelihood of support among liberals. A gun control framing, though, resulted in the opposite: conservatives become less likely to support gun ownership and liberals become more likely (Figure 5, left column). However, in the experimental (*word cloud* and *word web*) conditions, ideology had the same relationship with support for gun ownership, regardless of framing (Figure 5, center and right columns). Thus, the

Fig. 5. Interactions among framing, visualization condition, and political ideology in subsequent support for gun ownership (i.e., where no participant saw any visualization). As in Figure 4, a framing effect occurs in the Text Only condition; in the left column, slopes for Oppose and Support are opposite in the two framings. Also as in Figure 4, that framing effect does not occur in either of the visualization conditions; in the middle and right columns, relationship between ideology and subsequent support for gun ownership does not differ between the two frame conditions. Thus, the mitigating effect of the visualization can transfer to questions where the participant sees no visualization.

frame-mitigating effect of the visualization transferred to situations where the visualization was absent.

Of the six questions we examined, five questions evidenced a transfer effect: same-sex marriage [65], global climate change [74], free speech [71], gun ownership, and jobs [67]. The one question that did not show evidence for a transfer effect, the question on abortion, also exhibited no framing effect in the control (text only) condition. Thus, in every case where a framing effect was observed in the control condition, a transfer effect was observed with the visualizations.

5 DISCUSSION

This study examined the ways that visual and spatial methods for representing text may influence the cognitive and interpretive processes by which framing effects occur. Results show that the two text visualization variations tested have a significant effect on opinions elicited from participants. In most cases, these effects also transfer to future situations where no visualization is present. Furthermore, while prior work in political communication has noted that framing can interact with ideology [65, 74], our results demonstrate that text visualizations can mediate the nature of this interaction. This section first discusses explanations for, and interpretations of, these results. We then consider implications for the design and evaluation of tools for generating alternative representations of text in terms of reflective interpretative and cognitive processes.

5.1 Interpretation of Results

At least two explanations could account for the visualizations' ability to mitigate framing. First, framing effects occur via heuristic, largely subconscious processes [32, 35, 52]. Language is usually comprehended through reading words in a relatively linear fashion in the order in which they appear. Presenting the text through a visualization enables the viewer to consider alternate connotations and, in the case of word webs, explore connections not readily apparent in a linear reading. In this way, the text visualization may disrupt typical heuristic processes, switching participants into *more deliberate modes of cognition and interpretation* that are less susceptible to framing effects [25, 40, 41, 57, 89]. Second, alternatively, the visualization may simply *undermine or mask the frame*. In this case, participants may still have been thinking in a heuristic manner, but the visualization did not present language from the article in a way that activated a particular interpretive frame.

Two specific points support the deliberate cognition explanation. First, participants' open-ended responses about how they answered the opinion questions suggest that those in the visualization conditions did in fact approach both the issue and the questions about it in a manner that differed from those who were exposed to the Text Only view. For example, one participant used the word web to "highlight a word and look at the connections it had with the other words to try get a sense or make an assumption of what the article's stance or viewpoint is" (P#40, cost framing). Another participant "started by clicking the biggest words first and read the sections that were then highlighted in the corresponding article. I then looked at the smaller words connected to the larger ones and tried to make a connection between the two based on what I had just read" (P#110, equality framing). These responses attest to the non-linear interactions made possible by the word web visualization. In contrast, those in the Text Only condition, many responses resembled the following: "while answering questions regarding the author I took the text into account. I completely disregarded the text while answering questions that pertained to my opinion" (P#133, equality framing). One participant indicated that "the text made me think about both the state and the federal governments role in taking care of U.S. citizens. [...] However, I do not think that the text influenced me that much. I already had opinions about most of those topics" (P#109, cost framing). As noted above (final paragraph of Section 4.3.1), this stated reliance on one's own prior opinions demonstrates the subconscious nature of framing effects' influence [32, 35, 52].

Second, if the visualizations simply masked or undermined the frame, we would expect no transfer effects to occur. However, of the six questions we tested, all five in which a framing effect occurred in the control condition showed evidence of transfer effects. The presence of these transfer effects lends more credence to the explanation that the visualizations primed more deliberate cognitive processes that were less susceptible to framing effects.

From there, we offer two possible explanations for the transfer effects. First, the cognitive and interpretive modes primed [72] by the visualization may persist beyond the presence of the image.

In this case, it might not matter exactly what text was being visualized. Rather, the mere presence of a spatial representation of text may prime these alternative interpretive and cognitive modes. Second, a learning effect may occur. That is, after seeing the visualization, participants learn to read and interpret text in a particular way that undermines the effects of framing. In this case, actually seeing the text that was visualized plays a more important role, as it facilitates internalization of the relationships between prose and spatial presentations and the concomitant differences in emphasis and salience [30, 48]. One interesting question then becomes the duration of such learning effects. Do they last for minutes, hours, days, weeks, or longer? Both these explanations could be tested empirically, an important task for future work.

The astute reader may note that the Text Only condition with the cost framing leads to a somewhat counter-intuitive result: the most liberal respondents actually become *opposed* to national healthcare. It stands to reason that a cost-based framing might appeal to conservatives, who often emphasize the importance of economic issues [21]. How, though, could framing cause liberals to reverse their usual position on this issue? Druckman et al. [24], from whom the study materials were taken, do not provide data on interactions between ideology and framing effects. They do, however, provide a potential explanation. The Obama administration's strategic decision to avoid an inequality framing for the Affordable Care Act caused "some backlash among liberal supporters who felt betrayed by Obama's promise to work to lessen inequalities" [24, p. 448]. That said, significant future work, described further below, is necessary to understand in more detail how interactive text displays interact with such framing effects.

We expected that frame mitigation would only occur with the word web visualization. However, our results show that both the word web and word cloud mitigated framing effects, though the effect of the word cloud was only approaching statistical significance. It is possible that being exposed to any disruption of conventional linearity may sufficiently influence the cognitive processes by which framing effects occur. Further, the mere presence of the word cloud could have triggered participants to read text (any text) in a fairly deliberate, directed manner informed by visual conventions [48, 50]. As mentioned above, comments in the open-ended responses that describe using the word cloud as a summary suggest this as a distinct possibility. However, our study was not designed to probe links between specific interaction patterns and the cognitive and interpretive processes each visualization design afforded. Rather, the goal was to test whether varying impacts on framing mechanisms occurred, which they did.

5.2 The Evaluation of Text Visualization in Communicative Contexts

This study also contributes to a consistent stream of information visualization research that has called for expanding the techniques and frameworks for evaluating visual representations of data [9, 44, 53, 64, 75]. Our work responds to this call by evaluating interactive text visualizations not in terms of task performance but with respect to direct and indirect impacts on cognitive and interpretive processes. While prior work has also explored novel visualization techniques for legislative texts [3] or political opinion data [29], the focus remained largely task-oriented. The study described here draws attention to the role that visualizations play in cognitive activities related to the formulation of beliefs rather than the completion of specific tasks [61]. Doing so exposes important questions regarding how visualizations mediate, mitigate, disrupt, or otherwise interact with mainstream communicative practices. In concert with prior studies cited above, our research indicates the amount of work that remains to be done to expand available approaches for evaluation to correspond with the rich communicative contexts in which visualizations are now being deployed. We believe our method for assessing framing effects could be adapted to a range of interpretive tasks, thus pointing the way to a series of investigations to address these gaps in our current understanding.

Our findings show the value of developing new evaluation methods to accommodate communication-driven visualization techniques, such as word clouds, especially in cases where popular or vernacular use seems to defy more established empirically based assessment practices [78, 83]. As the means to create visual representations of data expand, audiences for visualizations will become increasingly diversified. Visualizations will be used in ways that system designers did not initially consider, possibly in ways that challenge long-held best practices within traditional visualization communities [39]. Rather than passing over the experiences of these vernacular users, or the communicative contexts in which they often deploy visualizations, we suggest that this is a fruitful area for exploring novel evaluation techniques. Better understanding of vernacular and communicative contexts of use can contribute to the creation of new classes of visualization systems to support a broader range of cognitive, social, and communicative activities.

5.3 Limitations and Future Work

One of the primary contributions of this work is to introduce the idea that visualizations can be evaluated in terms of both their direct and their indirect impacts on interpretive processes. While communication research has looked at the ability of certain types of images, such as violent photos, to *prompt* negative interpretations of news stories or political opinions [88], this study shows that visualizations can also be used to *mitigate* the effects of unconscious manipulations. However, these results emerge from the study of a small subset of issues using two similar types of text visualization. Future work should examine how visualizations can function as priming mechanisms with the ability to mitigate, or at least highlight and promote reflective awareness of, the mechanisms of other subconscious processes. For instance, behavioral economists have shown how perceptions of risk under uncertainty can differ based on whether the risk is described in terms of costs or benefits [67, 82]. Prior work has also identified a number of subtle effects that arise from demographic differences, varying from stereotype threat [80] to implicit racial bias when reviewing job applicants [10]. Our findings have implications for the design and evaluation of systems to help users understand the mechanisms of discrimination, prejudice, and many other sorts of biases that emerge from the presentation of data, such as census results; opinion polls; and education, race, and poverty demographics. In short, these results suggest an exciting and potentially important line of future work on text visualization and interactive display of information.

The results presented here provide evidence for the ability of text visualizations to influence cognitive and interpretive processes involved in reading text. However, this experiment was not designed to interrogate the mechanisms via which those influences occur. For instance, the system included no fine-grained logging, gaze tracking, or other quantitative data about specifically what participants did with each text and visualization. Since the three conditions being compared (Text Only, Word Cloud, and Word Web) did not include directly comparable features, the above study only collected self-report responses from participants describing their use of each text or visualization. Future work comparing different variants of the same interface would benefit from collecting such fine-grained behavioral data. Doing so would provide experimental means of assessing the impact of specific design elements not examined here, such as rotational orientation of the text, luminance or saturation values, the presence of overlapping words, and others.

Several aspects of our experimental design should also be noted. Although the experimental materials for both visualization conditions placed primacy on the visualization, they did so in a slightly different manner. The word cloud appeared above the text, while the word web appeared alongside it. Future work could examine differences between sequential vs. simultaneous presentation of visualization and text. Also, our random assignment of participants to experimental conditions resulted in ideological imbalance with respect to framing: conservative participants more often saw the Equality framing, and liberal participants more often saw the cost framing

(see Table 2). While this imbalance was unlikely to substantially impact the results, subsequent studies with more diverse and/or representative samples would help corroborate the results observed here. Finally, the experimental manipulation of framing was a simple either-or comparison where participants saw only a single frame. This simple experimental design was important as an initial step to demonstrate that an effect occurs. Subsequent research should draw on the political communication literature to explore more complex experimental designs, such as exposure to multiple frames, discussion among participants, or requiring participants to provide rationales for their decisions [24, 25, 40, 41, 57, 65, 89].

While we found evidence for a transfer effect, relatively little time passed between the initial and transfer effects. Subsequent work could consider how long such transfer effects last and in what contexts they apply. Also, prior work has questioned the ecological validity of framing studies that expose participants to a single frame [24]. More complex experimental designs would be able to probe the effects of text visualization under exposure to multiple competing frames of varying strength. Finally, because we wanted participants who were experts neither in political framing nor in text visualization, we recruited undergraduate students. Prior work has found that among some populations, such as Mechanical Turk workers, political ideology as a moderator functions differently than among a representative sample [51]. Future work should thus explore the impacts of text visualizations among other specific populations.

6 CONCLUSION

The visual and material form in which text is presented can significantly influence the cognitive and interpretive processes of reading text. Computational tools increasingly give designers the ability to apply consistent visual transformations to arbitrary text. However, we know relatively little about the various ramifications of such transformations.

This article studies one particular kind of transformation: disruption of the conventional linear structure of text. It examines the influence of replacing this display, either with a non-linear presentation (the word cloud) or with a spatial network structure (the word web). Using the framing of political issues as an example, the study presented here shows that both visual approaches result in the mitigation of previously documented framing effects associated with the conventional display of text. These findings demonstrate one of many ways that alternative visual representations can influence cognitive and interpretive processes. Furthermore, they also suggest how those influences occur: disrupting conventional linearity may circumvent the normal heuristic processes of reading text, resulting in more intentional, deliberate cognition, and interpretation.

Thus, from the contributions of this article, we draw three main implications. First, our results show that text visualization can effectively mitigate framing effects. This finding contributes important insights about the cognitive and interpretive processes by which framing effects occur and inspires novel approaches to the evaluation of visualizations in communicative contexts. Second, we suggest that alternative structural representations of text (including interactive and spatial displays) may play a central role in mitigating such subconscious processes. These findings provide important implications for exploring the design space of text visualization. Third, the frame-mitigating effects of visualizations can transfer to other situations where no visualization is present. These findings offer exciting possibilities for longitudinal assessment of the broader impacts of the visual display of information in a wide variety of contexts.

ACKNOWLEDGMENTS

We thank Jessica Hullman for comments on an earlier draft, and to our study participants. Portions of the work presented here were completed while the first author was affiliated with Communication and with Information Science at Cornell University.

REFERENCES

- [1] James Allen. 2002. *Topic Detection and Tracking: Event-Based Information Organization*. Kluwer Academic Publishers, New York.
- [2] Kevin Arceneaux. 2012. Cognitive biases and the strength of political arguments. *Am. J. Polit. Sci.* 56, 2 (April 2012), 271–285. DOI : <https://doi.org/10.1111/j.1540-5907.2011.00573.x>
- [3] Yannick Assogba, Irene Ros, Joan DiMicco, and Matt McKeon. 2011. Many bills: Engaging citizens through visualizations of congressional legislation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'11)*. 433–442.
- [4] David Bamman, Jacob Eisenstein, and Tyler Schnoebelen. 2014. Gender identity and lexical variation in social media. *J. Socioling* 18, 2 (2014), 135–160. DOI : <https://doi.org/10.1111/josl.12080>
- [5] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. 2009. Gephi: An open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM'09)*. 361–362.
- [6] Eric P. S. Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, and Geri K. Gay. 2015. Testing and comparing computational approaches for identifying the language of framing in political news. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics – Human Language Technologies (NAACL'15)*. 1472–1482.
- [7] Eric P. S. Baumer, Francesca Polletta, Nicole Pierski, and Geri K. Gay. 2017. A simple intervention to reduce framing effects in perceptions of global climate change. *Environ. Commun.* 11, 3 (2017), 289–310. DOI : <https://doi.org/10.1080/17524032.2015.1084015>
- [8] Victor Benjamin, Wingyan Chung, Ahmed Abbasi, Joshua Chuang, Catherine A. Larson, and Hsinchun Chen. 2013. Evaluating text visualization: An experiment in authorship analysis. In *Proceedings of the IEEE International Conference on Intelligence and Security Informatics (ISI'13)*. 16–20. DOI : <https://doi.org/10.1109/ISI.2013.6578778>
- [9] Enrico Bertini, Adam Perer, Catherine Plaisant, and Giuseppe Santucci. 2008. BELIV'08: Beyond time and errors: Novel evaluation methods for information visualization. In *Proceedings of the Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI EA'08)*. 3913–3916. DOI : <https://doi.org/10.1145/1358628.1358955>
- [10] Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than lakisha and jamal? A field experiment on labor market discrimination. *Am. Econ. Rev.* 94, 4 (2004), 991–1013. DOI : <https://doi.org/10.2307/3592802>
- [11] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O'Reilly Media, Inc., Sebastopol, CA.
- [12] David M. Blei. 2012. Probabilistic topic models. *Commun. ACM* 55, 4 (April 2012), 77–84. DOI : <https://doi.org/10.1145/2133806.2133826>
- [13] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 1, (2003), 993–1022.
- [14] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D3: Data-driven documents. *IEEE Trans. Vis. Comput. Graph. Proc. Info. Vis.* 17, 11 (2011), 2301–2309.
- [15] Eva Brumberger. 2004. The rhetoric of typography: Effects on reading time, reading comprehension, and perceptions of ethos. *Tech. Commun.* 51, 1 (2004), 13–24.
- [16] Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'15)*. 438–444.
- [17] Dennis Chong and James N. Druckman. 2007. Framing theory. *Annu. Rev. Polit. Sci.* 10, 1 (June 2007), 103–126. DOI : <https://doi.org/10.1146/annurev.polisci.10.072805.103054>
- [18] Jason Chuang, Daniel Ramage, Christopher D. Manning, and Jeffrey Heer. 2012. Interpretation and trust: Designing model-driven visualizations for text analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'12)*. 443–452. DOI : <https://doi.org/10.1145/2207676.2207738>
- [19] Weiwei Cui, Shixia Liu, Li Tan, Conglei Shi, and Yangqiu Song. 2011. TextFlow: Towards better understanding of evolving topics in text. *IEEE Trans. Vis. Comput. Graph.* 17, 12 (2011), 2412–2421. DOI : <https://doi.org/10.1109/TVCG.2011.239>
- [20] Evanthia Dimara, Anastasia Bezerianos, and Pierre Dragicevic. 2017. The attraction effect in information visualization. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (January 2017), 471–480. DOI : <https://doi.org/10.1109/TVCG.2016.2598594>
- [21] Michael Dimock, Jocelyn Kiley, Scott Keeter, Carroll Doherty, and Alec Tyson. 2014. *Beyond Red vs. Blue: The Political Typology*. Pew Research Center, Washington, D.C.
- [22] Marian Dörk, Patrick Feng, Christopher Collins, and Sheelagh Cpendale. 2013. Critical InfoVis: Exploring the politics of visualization. In *Proceedings of the Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI EA'13)*. 2189–2198. DOI : <https://doi.org/10.1145/2468356.2468739>

- [23] James N. Druckman. 2001. The implications of framing effects for citizen competence. *Polit. Behav.* 23, 3 (2001), 225–256.
- [24] James N. Druckman, Jordan Fein, and Thomas J. Leeper. 2012. A source of bias in public opinion stability. *Am. Polit. Sci. Rev.* 106, 2 (May 2012), 430–454. DOI: <https://doi.org/10.1017/S0003055412000123>
- [25] James N. Druckman and Kjersten R. Nelson. 2003. Framing and deliberation: How citizens' conversations limit elite influence. *Am. J. Polit. Sci.* 47, 4 (October 2003), 729–745. DOI: <https://doi.org/10.1111/1540-5907.00051>
- [26] Alex Endert, Patrick Fiaux, and Chris North. 2012. Semantic interaction for visual text analytics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'12)*. 473–482.
- [27] Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *J. Commun.* 43, 4 (1993), 51–58. DOI: <https://doi.org/10.1111/j.1460-2466.1993.tb01304.x>
- [28] Norman Fairclough. 1999. Global capitalism and critical awareness of language. *Lang. Aware.* 8, 2 (1999), 71–83.
- [29] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion Space: A scalable tool for browsing online comments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'10)*. 1175–1184. DOI: <https://doi.org/http://doi.acm.org/10.1145/1753326.1753502>
- [30] Stephen Few. 2009. *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. Analytics Press, Oakland, CA.
- [31] William A. Gamson. 1989. News as framing. *Am. Behav. Sci.* 33, 2 (November 1989), 157–161. DOI: <https://doi.org/10.1177/0002764289033002006>
- [32] William A. Gamson and Andre Modigliani. 1989. Media Discourse and public opinion on nuclear power: A constructionist approach. *Am. J. Sociol.* 95, 1 (July 1989), 1–37.
- [33] Benjamin Ginsberg. 1986. *The Captive Public: How Mass Opinion Promotes State Power*. Basic Books, New York.
- [34] Todd Gitlin. 1980. *The Whole World is Watching*. University of California Press, Berkeley, CA.
- [35] Erving Goffman. 1974. *Frame Analysis*. Harvard University Press, Cambridge, MA.
- [36] Frank van Ham, Martin Wattenberg, and Fernanda Viégas. 2009. Mapping text with phrase nets. *IEEE Trans. Vis. Comput. Graph.* 15, 6 (2009), 1169–1176. DOI: <https://doi.org/10.1109/TVCG.2009.165>
- [37] Marti A. Hearst and Daniela Rosner. 2008. Tag clouds: Data analysis tool or social signaller? In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS'08)*. 1–10.
- [38] Jeffrey Heer, Frank van Ham, Sheelagh Carpendale, Chris Weaver, and Petra Isenberg. 2008. Creation and collaboration: Engaging new audiences for information visualization. In *Information Visualization: Human-Centered Issues and Perspectives*, Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North (Eds.). Springer, Berlin, 92–133. DOI: https://doi.org/10.1007/978-3-540-70956-5_5
- [39] Jeff Hemsley and Jaime Snyder. 2018. Dimensions of visual misinformation in the emerging media landscape. In *Misinformation and Mass Audiences*, Biran G. Southwell, Emily A. Thorson, and Laura Sheble (Eds.). University of Texas Press, Austin, TX, 91–106.
- [40] Gerard P. Hodgkinson, Nicola J. Bown, A. John Maule, Keith W. Glaister, and Alan D. Pearman. 1999. Breaking the frame: An analysis of strategic cognition and decision making under uncertainty. *Strateg. Manag. J.* 20, 10 (1999), 977–985.
- [41] Gerard P. Hodgkinson, A. John Maule, Nicola J. Bown, Alan D. Pearman, and Keith W. Glaister. 2002. Further reflections on the elimination of framing bias in strategic decision making. *Strateg. Manag. J.* 23, 11 (November 2002), 1069–1076. DOI: <https://doi.org/10.1002/smj.266>
- [42] Jessica Hullman and Nick Diakopoulos. 2011. Visualization rhetoric: Framing effects in narrative visualization. *IEEE Trans. Vis. Comput. Graph.* 17, 12 (2011), 2231–2240. DOI: <https://doi.org/10.1109/TVCG.2011.255>
- [43] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. 2011. Collaborative visualization: Definition, challenges, and research agenda. *Inf. Vis.* 10, 4 (2011), 310–326. DOI: <https://doi.org/10.1177/1473871611412817>
- [44] Petra Isenberg, Torre Zuk, Christopher Collins, and Sheelagh Carpendale. 2008. Grounded evaluation of information visualizations. In *Proceedings of the Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV'08)*. 1–8. DOI: <https://doi.org/10.1145/1377966.1377974>
- [45] Shanto Iyengar. 1990. Framing responsibility for political issues: The case of poverty. *Polit. Behav.* 12, 1 (1990), 19–40.
- [46] Lawrence R. Jacobs. 2012. Barack Obama and the angry left. In *Obama at the Crossroads*, Lawrence R. Jacobs and Desmond S. King (Eds.). Oxford University Press, Oxford, 181–194.
- [47] Steven Johnson. 2013. Why no one clicked on the great hypertext story. *Wired*. Retrieved August 18, 2017 from <https://www.wired.com/2013/04/hypertext/>.
- [48] Charles Kostelnick and Michael Hassett. 2003. *Shaping Information: The Rhetoric of Visual Conventions*. Southern Illinois University Press, Carbondale, IL.
- [49] Gunther Kress and Theo Van Leeuwen. 2001. *Multimodal Discourse: The Modes and Media of Contemporary Communication*. Edward Arnold, London.

- [50] Gunther R. Kress and Theo Van Leeuwen. 1996. *Reading Images: The Grammar of Visual Design*. Psychology Press, London.
- [51] Yanna Krupnikov and Adam Seth Levine. 2014. Cross-sample comparisons and external validity. *J. Exp. Polit. Sci.* 1, 1 (August 2014), 59–80. DOI: <https://doi.org/10.1017/xps.2014.7>
- [52] George Lakoff and Mark Turner. 1989. *More Than Cool Reason: A Field Guide to Poetic Metaphor*. University of Chicago Press, Chicago and London.
- [53] Heidi Lam, Enrico Bertini, Petra Isenberg, Catherine Plaisant, and Sheelagh Carpendale. 2012. Empirical studies in information visualization: Seven scenarios. *IEEE Trans. Vis. Comput. Graph.* 18, 9 (September 2012), 1520–1536. DOI: <https://doi.org/10.1109/TVCG.2011.279>
- [54] Kevin Larson and Rosalind W. Picard. 2005. Measuring the aesthetics of reading. In *Human Computer Interaction Consortium*. DOI: https://doi.org/10.1007/978-1-84628-664-3_4
- [55] Jure Leskovec, Anand Rajaraman, and Jeffrey D. Ullman. 2014. Data mining. In *Mining of Massive Datasets*. Cambridge University Press, Cambridge, 1–19.
- [56] Neil Malhotra. 2008. Completion time and response order effects in web surveys. *Public Opin. Q.* 72, 5 (2008), 914–934. DOI: <https://doi.org/10.1093/poq/nfn050>
- [57] Paul M. Miller and N. S. Fagley. 1991. The effects of framing, problem variations, and providing rationale on choice. *Pers. Soc. Psychol. Bull.* 17, 5 (October 1991), 517–522. DOI: <https://doi.org/10.1177/0146167291175006>
- [58] Theodor Nelson. 1974. *Computer Lib*. Self-Published.
- [59] Theodor H. Nelson. 1965. Complex information processing: A file structure for the complex, the changing and the indeterminate. In *Proceedings of the National Conference of the ACM (ACM'65)*. 84–100. DOI: <https://doi.org/10.1145/800197.806036>
- [60] Sigrid Norris. 2004. *Analyzing Multimodal Interaction: A Methodological Framework*. Routledge, New York.
- [61] Chris North. 2006. Toward measuring visualization insight. *IEEE Comput. Graph. Appl.* 26, 3 (2006), 6–9. DOI: <https://doi.org/10.1109/MCG.2006.70>
- [62] Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Found. Trends® Inf. Retr.* 2, 1–2 (2008), 1–135. DOI: <https://doi.org/10.1561/15000000011>
- [63] Defen Peng and Gilbert MacKenzie. 2014. Discrepancy and choice of reference subclass in categorical regression models. In *Statistical Modelling in Biostatistics and Bioinformatics*, Gilbert MacKenzie and Defen Peng (Eds.). Springer International Publishing, 159–184. DOI: https://doi.org/10.1007/978-3-319-04579-5_12
- [64] Zachary Pousman, John T. Stasko, and Michael Mateas. 2007. Casual information visualization: Depictions of data in everyday life. *IEEE Trans. Vis. Comput. Graph.* 13, 6 (2007), 1145–1152. DOI: <https://doi.org/10.1109/TVCG.2007.70541>
- [65] Vincent Price, Lilach Nir, and Joseph N. Cappella. 2005. Framing public discussion of gay civil unions. *Public Opin. Q.* 69, 2 (2005), 179–212.
- [66] Harald E. L. Prins and John Bishop. 2002. Edmund carpenter: Explorations in media & anthropology. *Vis. Anthropol. Rev.* 17, 2 (2002), 110–140.
- [67] George A. Quattrone and Amos Tversky. 1988. Contrasting rational and psychological analyses of political choice. *Am. Polit. Sci. Rev.* 82, 3 (1988), 719–736.
- [68] Martin Rein and Donald A. Schön. 1996. Frame-critical policy analysis and frame-reflective policy practice. *Knowl. Policy* 9, 1 (1996), 85–104.
- [69] John Risch, Anne Kao, Stephen R. Poteet, and Y.-J. Jason Wu. 2008. Text visualization for visual text analytics. In *Visual Data Mining*, Simeon J. Simoff, Michael H. Böhlen, and Arturas Mazeika (Eds.). Springer, Berlin, 154–171. DOI: https://doi.org/10.1007/978-3-540-71080-6_11
- [70] Henry L. Roediger and Teresa A. Blaxton. 1987. Effects of varying modality, surface features, and retention interval on priming in word-fragment completion. *Mem. Cognit.* 15, 5 (1987), 379–388. DOI: <https://doi.org/10.3758/BF03197728>
- [71] Donald Rugg. 1941. Experiments in wording questions: II. *Public Opin. Q.* 5, 1 (1941), 91–92. DOI: <https://doi.org/10.1086/265467>
- [72] Dietram A. Scheufele and David Tewksbury. 2007. Framing, agenda setting, and priming: The evolution of three media effects models. *J. Commun.* 57, 1 (March 2007), 9–20. DOI: <https://doi.org/10.1111/j.1460-2466.2006.00326.x>
- [73] Donald A. Schön and Martin Rein. 1994. *Frame Reflection: Toward the Resolution of Intractable Policy Controversies*. Basic Books, New York.
- [74] Jonathon P. Scholdt, S. H. Konrath, and N. Schwarz. 2011. “Global warming” or “climate change”? Whether the planet is warming depends on question wording. *Public Opin. Q.* 75, 1 (February 2011), 115–124. DOI: <https://doi.org/10.1093/poq/nfq073>
- [75] Michael Sedlmair, Petra Isenberg, Tobias Isenberg, Narges Mahyar, and Heidi Lam. 2016. *Proceedings of the Sixth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization (BELIV 2016)*. ACM.
- [76] Artur Šilić and Bojana Dalbelo Bašić. 2010. Visualization of text streams: A survey. In *Proceedings of the Knowledge-Based and Intelligent Information and Engineering Systems (KES'10)*. 31–43.

- [77] Paul M. Sniderman and Sean M. Theriault. 2004. The structure of political argument and the logic of issue framing. In *Studies in Public Opinion*, W. E. Saris and Paul M. Sniderman (Eds.). Princeton University Press, Princeton, NJ, 133–165.
- [78] Jaime Snyder. 2017. Vernacular visualization practices in a citizen science project. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW'17)*. 2097–2111. DOI: <https://doi.org/10.1145/2998181.2998239>
- [79] Jamie Snyder, Michael A. D'Eredita, Ozgu Yilmazel, and Elizabeth D. Liddy. 2009. Towards a cognitive approach for the automated detection of connotative meaning. In *Proceedings of the International Conference on Computational Semantics (IWCS'09)*. 326–332.
- [80] Claude M. Steele and Joshua Aronson. 1995. Stereotype threat and the intellectual test performance of African Americans. *J. Pers. Soc. Psychol.* 69, 5 (1995), 797–811. DOI: <https://doi.org/10.1037/0022-3514.69.5.797>
- [81] Teal Triggs. 2003. *The Typographic Experiment: Radical Innovation in Contemporary Type Design*. Thames & Hudson, London.
- [82] Amos Tversky and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211, 4481 (January 1981), 453–458.
- [83] Fernanda B. Viégas and Martin Wattenberg. 2008. Tag clouds and the case for vernacular visualization. *Interactions* 15, 4 (2008), 49. DOI: <https://doi.org/10.1145/1374489.1374501>
- [84] Justin Washtell and Katja Markert. 2009. A comparison of windowless and window-based computational association measures as predictors of syntagmatic human associations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'09)*. 628–637.
- [85] Martin Wattenberg and Fernanda B. Viégas. 2008. The word tree, an interactive visual concordance. *IEEE Trans. Vis. Comput. Graph.* 14, 6 (2008), 1221–1228. DOI: <https://doi.org/10.1109/TVCG.2008.172>
- [86] Furu Wei, Shixia Liu, Yangqiu Song, Shimei Pan, Michelle Zhou X., Weihong Qian, Lei Shi, Li Tan, and Qiang Zhang. 2010. TIARA: A visual exploratory text analytic system. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'10)*. 153–162. DOI: <https://doi.org/10.1145/1835804.1835827>
- [87] James A. Wise. 1999. The ecological approach to text visualization. *J. Am. Soc. Inf. Sci. Technol. JASIST* 50, 13 (1999), 1224–1233. DOI: [https://doi.org/10.1002/\(SICI\)1097-4571\(1999\)50:13<1224::AID-ASI8>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1097-4571(1999)50:13<1224::AID-ASI8>3.0.CO;2-4)
- [88] Bernd Wittenbrink. 2007. Measuring attitudes through priming. In *Implicit Measures of Attitudes*, Bernd Wittenbrink and Norbert Schwarz (Eds.). The Guilford Press, New York, 17–58.
- [89] George Wright and Paul Goodwin. 2002. Eliminating a framing bias by using simple instructions to “think harder” and respondents with managerial experience: Comment on “breaking the frame”. *Strateg. Manag. J.* 23, 11 (November 2002), 1059–1067. DOI: <https://doi.org/10.1002/smj.265>
- [90] Koji Yatani, Michael Novati, Andrew Trusty, and Khai N. Truong. 2011. Review spotlight. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'11)*. 1541–1550. DOI: <https://doi.org/10.1145/1978942.1979167>
- [91] Ji Soo Yi, Youn-ah Kang, John T. Stasko, and Julie Jacko. 2007. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans. Vis. Comput. Graph.* 13, 6 (2007), 1224–1231.
- [92] Yunfeng Zhang, Rachel K. E. Bellamy, and Wendy A. Kellogg. 2015. Designing information for remediating cognitive biases in decision-making. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'15)*. 2211–2220. DOI: <https://doi.org/10.1145/2702123.2702239>
- [93] Jian Zhao, Fanny Chevalier, Christopher Collins, and Ravin Balakrishnan. 2012. Facilitating discourse analysis with interactive visualization. *IEEE Trans. Vis. Comput. Graph.* 18, 12 (2012), 2639–2648.

Received November 2016; revised August 2017; accepted March 2018