

Narrative Sensemaking: Strategies for Narrative Maps Construction

Brian Keith Norambuena*
Virginia Tech
Universidad Católica del Norte

Tanushree Mitra†
University of Washington

Chris North‡
Virginia Tech

ABSTRACT

Narrative sensemaking is a fundamental process to understand sequential information. Narrative maps are a visual representation framework that can aid analysts in this process. They allow analysts to understand the big picture of a narrative, uncover new relationships between events, and model connections between storylines. As a sensemaking tool, narrative maps have applications in intelligence analysis, misinformation modeling, and computational journalism. In this work, we seek to understand how analysts construct narrative maps in order to improve narrative map representation and extraction methods. We perform an experiment with a data set of news articles. Our main contribution is an analysis of how analysts construct narrative maps. The insights extracted from our study can be used to design narrative map visualizations, extraction algorithms, and visual analytics tools to support the sensemaking process.

Index Terms: Human-centered computing—Visualization—Empirical studies in visualization; Human-centered computing—Visualization—Visualization techniques—Graph drawings

1 INTRODUCTION

Narratives are systems of stories [13]—sequences of events tied together in a coherent fashion. Events are the fundamental units of narrative action, they are either an act involving characters and entities or a happening where no entities are causally involved [1]. Narratives are fundamental to our understanding of the world and provide a natural way to capture relationships between sequences of events, as well as the goals, motivations, and plans of actors [10]. Narratives are used in the process of “connecting the dots” between apparently unrelated pieces of information [15, 16].

Storytelling in general is an accepted metaphor used in visual analytics and analytical reasoning [7, 27, 31]. However, unlike general visual storytelling, our work focuses specifically on visualizing textual narratives, such as those created by news. In this context, narratives provide a way to understand the information landscape, a key part of several narrative sensemaking tasks [21]. For example, aiding in tasks such as journalistic analysis of news narratives [5] and intelligence analysis [9]. To aid analysts with sensemaking tasks, scholars have created visual analytics software, which allow analysts to process and understand greater quantities of data and information [7]. These tools focus on different parts of the sensemaking loop [24], such as foraging [20] and synthesis [33].

In this work, we focus on a specific type of graph-based narrative representation—narrative maps [21]. Narrative maps provide a generic foundation to encode any narrative extracted from data, requiring only the existence of a total ordering (e.g., timestamps) and text representation of the event (e.g., news headlines). Narrative maps are a useful visualization framework to understand the information landscape. As a sensemaking tool, they have applications in intelligence analysis, misinformation modeling, and computational

journalism [21]. In particular, they offer a way to keep track of the big picture of a narrative in the context of the ever-increasing problem of information overload [14, 28]. However, from a visualization standpoint, the construction process of narrative maps for sensemaking remains unexplored. Thus, our goal is to understand how analysts construct narrative maps to design better narrative representation and extraction methods. We focus on the strategies, cognitive connections, and structures used in the construction process.

2 RELATED WORK

Narratives are systems of stories with coherent themes [13]. Narrative studies attempt to understand the relationships between the underlying stories and their representations [1, 25]. In the context of information visualization, we explore how information narratives can be visualized. The narrative visualization pipeline requires extracting the narrative from data into an internal representation, which is then used to generate the final visualization [21].

Storytelling and narratives are common metaphors in visual analytics [31]. In general, scholars have studied how arranging visualizations as story sequences helps in sensemaking [18, 19]. Other works focused on narrative visualization for news usually focus on augmenting data visualization techniques (e.g., charts) with contextual information (e.g., relevant articles associated with data points in the chart) [11, 17]. However, in our application context, we are interested in extracting and representing narratives taken directly from data sets of text documents, rather than augmenting numerical (or other non-text types of data) visualizations with contextual information or using sequences of visualizations to represent a story.

Previous research has explored how analysts make cognitive connections between documents in the context of intelligence analysis tasks. Bradel et al. [4] studied how analysts structure information, finding layouts based on linear structures with branching and web-like structures. Robinson [26] focused on analyzing the strategies and organizational methods used in collaborative synthesis, with the purpose of proposing design guidelines for collaborative sensemaking systems. Andrews et al. [2] explored the workspace organization used by analysts in large displays to arrange documents, where most strategies consisted of clustering, although some analysts used timelines. Wenskovitch and North [32] studied how analysts perform grouping and dimensionality reduction, where strategies included divide-and-conquer, incremental layouts, and bottom-up construction. Our work follows a similar approach, but focusing exclusively on sensemaking with narrative maps, analyzing the different map construction strategies and structures generated in the process.

Previous studies have found that analysts use strategies such as identifying co-occurrence relationships and aggregating common elements [12], using topical and temporal orderings for document clustering, and evaluating content overlap and similarity for document summarization [6, 8]. However, previous research has not focused on narratives, which have an underlying temporal ordering as well as a focus on cause-effect relationships that leads to a specific description of cognitive connections and construction strategies.

Finally, prior works have shown that graph-based narrative representations [21, 22, 30] are useful as a sensemaking tool. Thus, with the purpose of improving such narrative representations and their extraction algorithms, we seek to understand how analysts create such models from scratch by analyzing the narrative mapping process.

*e-mail: briankeithn@vt.edu

†e-mail: tmitra@uw.edu

‡e-mail: north@vt.edu

3 STUDY DESCRIPTION

3.1 Data Set

We used a data set comprised of 40 COVID-19 news articles from January 2020 covering the start of the Coronavirus outbreak. This data set is a subset of the COVID-19 archive data used in previous works on narrative maps [21]. The events were carefully curated in order to have a sufficiently small data set while covering a series of different topics and issues regarding the COVID-19 narrative. In particular, the articles cover topics such as the economic consequences of the pandemic, the sociopolitical effects in China, and the worldwide response. As our data set was made up of breaking news, the main event is usually described explicitly in the headlines [23]. Thus, we focused on the headlines rather than the full article. We also included the publication dates and sources.

3.2 Task Definitions

We defined two tasks to explore how analysts constructed narrative maps, a directed task that required participants to join two events and an open-ended task that required participants to expand on the outcomes of an initial event. In both tasks, participants were given a list of events (i.e., nodes) and asked to construct a narrative map by designing its overall structure, layout, and specific connections. The participants were also asked to label their main storyline—the core events of the narrative—and their side stories—stories relevant to the overall narrative but not directly related to the main storyline.

The directed task required participants to construct a narrative map to answer the following question: *“How did the Wuhan outbreak lead to the US travel restrictions?”*, which referred to two specific events in the data set. This task is also known as “connecting the dots” and it is a fundamental task in narrative sensemaking [21]. Previous research has attempted to understand how analysts perform this process [4] and sought to automate this process through algorithmic approaches [28]. In contrast, the open-ended required participants to construct a narrative map to answer the following question: *“What outcomes occurred as a result of the Wuhan outbreak?”*. This task is a variation of the “connect the dots” task [29] that only provides the starting event as a fixed point, requiring participants to explore the stories that emerge because of this event.

Finally, both tasks required participants to label their storylines and to answer a follow-up question with their map: *“What are the key events (i.e., the most important events or turning points)?”*. The focus of this experiment was to glean insights on the construction process, rather than comparing how the tasks themselves influence the construction. By considering two tasks rather than a single one, we expected to gather additional insights on this process.

3.3 Evaluation Procedure

We recruited ten participants, following a similar approach to the work of Bradel et al. [4]. We assigned five participants to each task. While splitting the participants into two tasks increases variability, we expected to gather a wider range of construction strategies by doing this. All participants were part of a national security program and hence, had a background in intelligence analysis. They also had previous knowledge on the topic which they could leverage while conducting the tasks. Prior knowledge ranged from general knowledge about COVID-19 to stronger backgrounds since some participants were ardent followers of the pandemic news right from its start. Figure 1 shows examples of the maps created in each task.

To provide initial training and to avoid inducing biases in subsequent task performance, participants were provided with a short example narrative map on a different topic. We engaged with our participants in an hour-long semi-structured session in a video call where they completed their assigned narrative sensemaking task. We encouraged the participants to think aloud and ask questions and share any observations as they worked. All participants were

recorded and the videos were analyzed to understand their construction strategies. We explained that there were no correct or incorrect answers; as our goal was to understand the cognitive strategies used by the analysts to complete the tasks. However, the quality and conceptual cohesion varied among the solutions.

To construct the map, we gave participants a canvas on Google Drawings with the instructions and the list of articles chronologically ordered. The participants had to drag and drop the articles into the available space. Then, they had to add connections with arrows. The participants were instructed to design the map with other analysts as potential users in mind. The participants were familiar with Google Drawings and similar editing tools, thus they did not require additional training in its use, even if it might not have been their preferred tool for such an exercise. Moreover, they had full access to this tool through their institutional accounts.

We opted for Google Drawings in our study for several reasons. First, it provided a closer approximation of what a computational narrative map tool would look like compared to an approach using hand-drawn notes. Thus, even though it might influence the kinds of strategies used by the participants, these strategies should be closer to what we would expect with a computational tool. Second, given the limitations caused by the pandemic, using Google Drawings allowed us to do virtual sessions, thus minimizing the risks for the participants. Finally, it also provided a detailed editing history which, in conjunction with the recorded sessions, was useful to precisely analyze the steps taken by the participants.

4 NARRATIVE MAP CONSTRUCTION STRATEGIES

We now present our results, including the cognitive connections, the construction strategies, and the graph and layout properties of the maps. Throughout this process, we used open coding to discover the different types of cognitive strategies and analyze the results.

4.1 How do analysts connect events?

To answer this question, we asked participants to explain their connection strategies as they constructed the map as well as in the follow-up interviews. We identified seven types of connections, which we further divided into low-level, high-level, and supporting connections. Low-level connections are those that can be made directly from the content of the document (e.g., dates, keywords, entities present) without an in-depth analysis. In contrast, high-level connections involve applying cognitive schemas to synthesize information between events [4]. Supporting connections are used in conjunction with high-level connections as an auxiliary strategy to help connect events. For example, a connection could be based on cause-effect relationships between events (high-level connection) and speculation (supporting connection). These results show that beyond offering a chronological chain of events, it would be beneficial to visualize narratives based on different types of relationships, such as those based on similarity, entities, topics, and causality.

4.1.1 Low-level Connections

We identified 3 low-level connections (temporal, similarity, and entity) based on our analysis of the sessions with the participants.

1. **Temporal connections** (A happened before B): Most participants used temporal connections as their default strategy. In particular, this type of connection was used when there were no other explicit relationships between events and participants wanted to maintain the temporal sequence of the events (e.g., “I just followed chronological order”). All participants used temporal connections in their maps, because of the inherent chronology of narratives.

2. **Similarity connections** (A is similar to B): Users determined two events as similar primarily based on keyword matching and a superficial similarity evaluation (e.g., “All these events mentioned markets or production”).

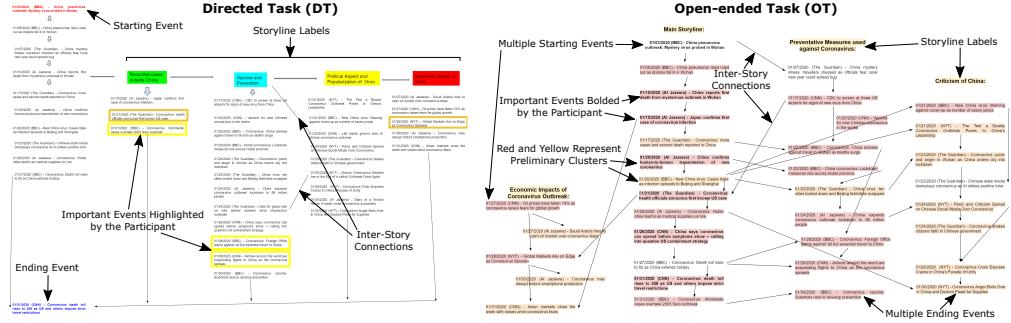


Figure 1: Narrative map examples created by participants for the two tasks **Directed Task (DT)** and **Open-ended Task (OT)**. Annotations highlight key elements in these maps. Note how in the DT map all storylines converge into a single ending event, emphasizing the focused nature of this task. In contrast, the OT map has a series of storylines that interact with one another, representing the different outcomes found by the participant.

3. Entity connections (A is about the same entity as B): These connections are based on named entity co-occurrences in events. For example, some participants focused on whether the events referred to specific entities; in this case, the entities mostly referred to the countries being affected (e.g., “These events talk about China”).

4.1.2 High-level Connections

We classified connections as high-level if they involved the use of a cognitive schema to connect information between documents. In particular, these connections usually arise from inferences made by the users rather than a superficial characteristic of the document.

1. Topical connections (A shares a common theme or topic with B): These connections are a more abstract version of the similarity connection. They focus on the overarching topic or theme of the articles (e.g., “These events are about the Chinese government response”). They differ from their low-level counterpart because they are based on a semantic viewpoint rather than superficial similarity.

2. Causal connections (A leads to B): These relationships join events if one is caused (or could be caused) by another (e.g., “The number of cases surpassing SARS led to stricter travel restrictions”). Causal connections also cover events that could be reactions to another event, even if they are not explicitly caused by it. Some causal relationships defy the temporal order because the reporting date is not the actual event date, leading to participants changing the order of the documents to respect the cause-effect relation.

4.1.3 Supporting Connections

We classified connections as supporting if they are auxiliary strategies used in conjunction with a high-level connection.

1. Speculative connections (A is connected to B because of inferred X): These connections are based on implications based on the participant’s beliefs (e.g., “The call for a global ban on animal markets made the global markets panic”, a paraphrased causal speculative connection from one of the participants). Thus, speculative connections relate events that do not share any explicit relationship but could be connected based on speculative reasons.

2. Domain Knowledge connections (A is related to B because of external knowledge X): These connections are a special type of connection where documents that do not share any explicit relationship are connected based on external domain knowledge (e.g., “Air travel and oil demand are related”).

4.2 What are analysts’ map construction strategies?

We studied the construction process by following the individual steps taken by the participants as they built their narrative maps. We also asked follow-up questions about the process during the interviews. The identified strategies are summarized in Figure 2.

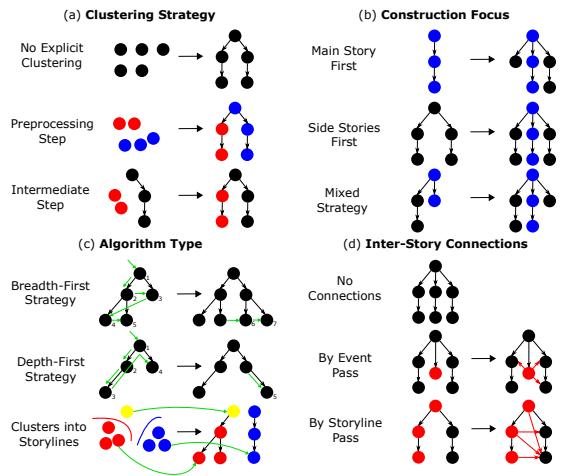


Figure 2: Narrative map construction strategies. (a) **Clustering strategies**: no clustering, clustering during preprocessing, or clustering in the middle of construction. (b) **Construction focus**: whether participants created the main story (in blue) first, the side stories first, or a mixed strategy. (c) **Algorithm type**: the order in which nodes were added, following either a breadth-first, depth-first, or clustering approach. (d) **Inter-story connections**: some participants checked for inter-story connections (in red) when adding events, others checked when completing a storyline.

1. Clustering Strategy: Clustering allows analysts to group documents based on specific characteristics (e.g., topic, type of document, source). Half of the participants had an explicit clustering step during the creation of the map. The use of clustering in sensemaking tasks has also been reported in previous research, either as a story construction strategy [4] or as the final product [8]. However, in the context of narrative maps, the main purpose of clustering is as a tool to aid in storyline constructions. If the participants performed clustering, it was either done as a preprocessing step or as an intermediate step (i.e., after starting with the connections). We note that clusters are not necessarily the same as storylines, but they can be used as a stepping stone towards identifying the storylines.

2. Construction Focus: This strategy refers to the part of the narrative map that was created first. Participants either focused on the main story, the side stories, or followed no particular order (i.e., a mixed strategy going back and forth). The main story refers to the sequences of core events in the narrative, those that move the narrative forward [1, 21]. In contrast, the side stories do not form part of the narrative core. Instead, they provide further information and useful context to the narrative. Note that there is an even split between focusing on the main story and following a mixed strategy. Only two participants decided to focus on the side stories first.

3. Algorithm Type: This strategy refers to the general algorithm that participants followed to construct the map. By analyzing the order in which participants constructed the maps, we found three types of strategies. The first two strategies are conceptually similar to basic graph searching algorithms—constructing the map in a depth-first or breadth-first fashion—while the third strategy is based on clustering—turning clusters into storylines. Note that depth-first approaches either focused on side stories first or on the main story first. In contrast, breadth-first approaches followed a mixed strategy by definition. The strategy of turning clusters into storylines either focused on side stories first or followed a mixed strategy.

4. Inter-story Connections: This strategy refers to how the participants connected storylines. In most cases, participants did not add inter-story connections; making their storylines independent from the rest of the graph, except for the initial connection where they split off. In other cases, they added connections on a by-event basis, checking whether an event should be connected to other stories as they add it. Alternatively, they added connections on a by-storyline basis, checking whether to connect the storyline with others only after completing the whole storyline.

4.3 What are the properties of the created maps?

We focus on multiple structural aspects of the underlying graph and the layout considerations made by participants (see Table 1).

Property	Code	DT	OT	Total
Graph Structure	List	1	1	2
	Tree	2	2	4
	DAG	2	2	4
Layout	Vertical (top-down)	3	4	7
	Diagonal (left to right)	1	0	1
	Horizontal (left to right)	1	1	2
Main Story Position	Main Story First	4	2	6
	Main Story Center	1	3	4
Source Nodes	Single	4	2	6
	Multiple	1	3	4
Sink Nodes	Single	2	0	2
	Multiple	3	5	8
Connectivity	Connected	4	4	8
	Disconnected	1	1	2
Transitivity	Implicit	5	3	8
	Explicit	0	2	2

Table 1: Graph and Layout Properties for each participant in our study.

1. Graph Structure: We found that participants used three types of underlying graph structures: lists , trees , and directed acyclic graphs (DAGs) .

These results are in line with prior work on narrative representations, which has focused on similar structures to represent stories [21], such as timelines [28], trees [3], or other graph variants [34]. Structures were evenly split between trees and DAGs, with only two list-like graphs, where one of them was a single timeline and the other comprised three parallel timelines .

2. Layout and Main Story Position: Most participants went for a vertical (top-down) approach with storylines presented as parallel columns and the main story placed first (i.e., the left-most story in a vertical layout or the top story in a horizontal layout). Horizontal layouts were not preferred; as noted by participants, computer displays seem to favor vertical layouts due to how scrolling works. Finally, one participant used a unique diagonal layout ; we did not observe this behavior in any of the other participants.

3. Number of Source and Sink Nodes: Most people had multiple storyline endings (i.e., sink nodes) . In particular, all five open-ended task maps had multiple endings. In contrast, the directed task had two participants constrain themselves to a single ending as defined by the task , while the others added endings or dead-end

events for some of the other storylines. For source nodes, participants that had the directed task were more likely to have a single source than those that had the open-ended task .

The tendency of open-ended maps to have multiple sources and sinks intuitively makes sense given the unrestricted nature of the task. In contrast, the directed task maps are focused on just answering the main question, thus leading to structures that did not have as many loose ends.

4. Connectivity: In graph-theoretical terms, we classify the map as connected if its underlying graph is weakly connected (i.e., we disregard the direction of the arrows). Most participants created connected graphs , but there were two cases with separate components .

The first had a separate component for the “social response and effects of COVID”. The second had three parallel timeline structures without any explicit connection between them.

5. Transitivity: We considered whether participants explicitly included connections that are implied by transitivity (i.e., vs.). Most people did not include these transitive connections. Only two participants who worked on the open-ended task used explicit transitive connections to emphasize the relationship between events. However, even in the maps where they were used, they were scarce. Thus, in general, transitive connections were either not needed or participants had difficulty finding them in the first place.

5 LIMITATIONS

Our work is not without limitations. In particular, we conducted interviews with only a handful of participants (10). While the number was small, all participants had a background in intelligence analysis. They also spanned a variety of majors and had reasonable gender representation (6 females and 4 males). Participants had different levels of expertise on the topic and were able to bring insights from their own knowledge and experiences. We note that experience and prior knowledge might influence the strategies, especially due to the use of a high-profile topic. Moreover, our data is strictly focused on COVID-19, which may bias participants to adopt specific strategies. However, even with these limitations, we were able to observe diverse strategies and structures for narrative map construction. Furthermore, we note that both of the tasks used in this study represent simplified and constrained versions of what analysts would do in a real-world setup, but they still provide valuable insights into the general narrative sensemaking process. Nevertheless, these limitations should be considered when attempting to generalize these conclusions, especially as higher complexity tasks or data sets might yield different kinds of strategies or structures.

6 CONCLUSIONS

We studied how analysts construct narrative maps and the characteristics of these maps. In particular, our user study detected 7 types of cognitive connections. We have shown the importance of topical and causal relationships in the construction of narrative maps, as these were the most common high-level connections in the user-generated maps. In terms of strategies, we found three major ways to construct maps. Each one of these strategies can be the basis of a new narrative extraction algorithm. Furthermore, in terms of the structure of the map, we saw an even distribution between tree-like maps and DAG-like maps. Regarding layout, we found that most users preferred a vertical top-down layout, with the main story shown first. These results can be used to define a series of design guidelines for narrative maps, as well as guide extraction algorithms and interactive visualization tools. Thus, future work should focus on developing specific design guidelines based on these results. Moreover, it would also be useful to explore how strategies differ when applied to different domains, data set sizes, and analyst experience. In particular, it would be useful to consider how previous analyst training (e.g., experience with structured analytic techniques) could influence the construction strategies or the narrative map structures.

ACKNOWLEDGMENTS

The authors wish to thank A, B, and C. This work was supported in part by a grant from XYZ.

REFERENCES

- [1] H. P. Abbott. *The Cambridge introduction to narrative*. Cambridge University Press, One Liberty Plaza, New York, NY, USA, 2008.
- [2] 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*, pp. 55–64, 2010.
- [3] J. Ansah, L. Liu, W. Kang, S. Kwashie, J. Li, and J. Li. A graph is worth a thousand words: Telling event stories using timeline summarization graphs. In *The World Wide Web Conference, WWW '19*, p. 2565–2571. ACM, New York, NY, USA, 2019.
- [4] L. Bradel, J. Z. Self, A. Endert, M. S. Hossain, C. North, and N. Ramakrishnan. How analysts cognitively “connect the dots”. In *2013 IEEE International Conference on Intelligence and Security Informatics*, pp. 24–26. IEEE, New York, NY, USA, 2013. doi: 10.1109/ISI.2013.6578780
- [5] L. Bradel, N. Wycoff, L. House, and C. North. Big text visual analytics in sensemaking. In *2015 Big Data Visual Analytics (BDVA)*, pp. 1–8. IEEE, 2015.
- [6] R. T. Camargo, V. Agostini, A. Di Felippo, and T. A. Pardo. Manual typification of source texts and multi-document summaries alignments. *Procedia-Social and Behavioral Sciences*, 95:498–506, 2013.
- [7] K. A. Cook and J. J. Thomas. Illuminating the path: The research and development agenda for visual analytics. Technical report, Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2005.
- [8] A. Endert, S. Fox, D. Maiti, S. Leman, and C. North. The semantics of clustering: Analysis of user-generated spatializations of text documents. In *Proceedings of the International Working Conference on Advanced Visual Interfaces, AVI '12*, p. 555–562. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2254556.2254660
- [9] A. Endert, M. S. Hossain, N. Ramakrishnan, C. North, P. Fiaux, and C. Andrews. The human is the loop: new directions for visual analytics. *Journal of intelligent information systems*, 43(3):411–435, 2014.
- [10] M. A. Finlayson and S. R. Corman. The military interest in narrative. *Sprache und Datenverarbeitung*, 37(1-2):173–191, 2013.
- [11] T. Gao, J. R. Hullman, E. Adar, B. Hecht, and N. Diakopoulos. Newsviews: an automated pipeline for creating custom geovisualizations for news. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 3005–3014, 2014.
- [12] J. D. Haider, P. Seidler, M. Pohl, N. Kodagoda, R. Adderley, and B. W. Wong. How analysts think: sense-making strategies in the analysis of temporal evolution and criminal network structures and activities. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 61, pp. 193–197. SAGE Publications: Los Angeles, CA, 2017.
- [13] J. Halverson, S. Corman, and H. L. Goodall. *Master narratives of Islamist extremism*. Springer, 175 5th Ave., New York, NY, USA, 2011.
- [14] J. Ho and R. Tang. Towards an optimal resolution to information overload: An infomediary approach. In *Proc. of the 2001 Int. ACM SIGGROUP Conference on Supporting Group Work, GROUP '01*, p. 91–96. ACM, New York, NY, USA, 2001.
- [15] M. S. Hossain, C. Andrews, N. Ramakrishnan, and C. North. Helping intelligence analysts make connections. In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.
- [16] M. S. Hossain, J. Gresock, Y. Edmonds, R. Helm, M. Potts, and N. Ramakrishnan. Connecting the dots between pubmed abstracts. *PloS one*, 7(1):e29509, 2012.
- [17] J. Hullman, N. Diakopoulos, and E. Adar. Contextifier: automatic generation of annotated stock visualizations. In *Proceedings of the SIGCHI Conference on human factors in computing systems*, pp. 2707–2716, 2013.
- [18] J. Hullman, S. Drucker, N. H. Riche, B. Lee, D. Fisher, and E. Adar. A deeper understanding of sequence in narrative visualization. *IEEE Transactions on visualization and computer graphics*, 19(12):2406–2415, 2013.
- [19] J. Hullman, R. Kosara, and H. Lam. Finding a clear path: Structuring strategies for visualization sequences. In *Computer Graphics Forum*, vol. 36, pp. 365–375. Wiley Online Library, 2017.
- [20] Y. Kang, C. Gorg, and J. Stasko. Evaluating visual analytics systems for investigative analysis: Deriving design principles from a case study. In *2009 IEEE Symposium on Visual Analytics Science and Technology*, pp. 139–146. IEEE, New York, NY, USA, 2009. doi: 10.1109/VAST.2009.5333878
- [21] B. Keith Norambuena and T. Mitra. Narrative maps: An algorithmic approach to represent and extract information narratives. In *Proc. ACM Hum.-Comput. Interact.*, vol. 4, p. 33 pages. ACM, New York, NY, USA, December 2020.
- [22] B. Liu, D. Niu, K. Lai, L. Kong, and Y. Xu. Growing story forest online from massive breaking news. In *Proc. of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, p. 777–785. ACM, New York, NY, USA, 2017.
- [23] B. K. Norambuena, M. Horning, and T. Mitra. Evaluating the inverted pyramid structure through automatic 5w1h extraction and summarization. In *Proc. of the 2020 Computation + Journalism Symposium*, pp. 1–7. Computation + Journalism 2020, Boston, MA, USA, 2020.
- [24] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4. International Conference on Intelligence Analysis, McLean, VA, USA, 2005.
- [25] K. Puckett. *Narrative theory*. Cambridge University Press, One Liberty Plaza, New York, NY, USA, 2016.
- [26] A. C. Robinson. Collaborative synthesis of visual analytic results. In *2008 IEEE Symposium on Visual Analytics Science and Technology*, pp. 67–74. IEEE, 2008.
- [27] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE transactions on visualization and computer graphics*, 16(6):1139–1148, 2010.
- [28] D. Shahaf and C. Guestrin. Connecting the dots between news articles. In *Proc. of the 16th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, KDD '10*, p. 623–632. ACM, New York, NY, USA, 2010.
- [29] D. Shahaf and C. Guestrin. Connecting two (or less) dots: Discovering structure in news articles. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):1–31, 2012.
- [30] D. Shahaf, C. Guestrin, and E. Horvitz. Metro maps of information. *SIGWEB Newslett.*, Spring, 2013.
- [31] C. Tong, R. Roberts, R. Borgo, S. Walton, R. S. Laramee, K. Wegba, A. Lu, Y. Wang, H. Qu, Q. Luo, et al. Storytelling and visualization: An extended survey. *Information*, 9(3):65, 2018.
- [32] J. Wenskovitch and C. North. An examination of grouping and spatial organization tasks for high-dimensional data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [33] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: Concepts and methods. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06*, p. 801–810. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1124772.1124890
- [34] C. C. Yang, X. Shi, and C.-P. Wei. Discovering event evolution graphs from news corpora. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 39(4):850–863, 2009.