Literal Encoding: Text is a first-class data encoding

Richard Brath*
Uncharted Software Inc.

arXiv:2009.02374v1 [cs.HC] 4 Sep 2020

ABSTRACT

Digital humanities are rooted in text analysis. However, most visualization paradigms use only categoric, ordered or quantitative data. Literal text must be considered a base data type to encode into visualizations. Literal text offers functional, perceptual, cognitive, semantic and operational benefits. These are briefly illustrated with a subset of sample visualizations focused on semantic word sequences, indicating benefits over standard graphs, maps, treemaps, bar charts and narrative layouts.

Index Terms: Human-centered computing—Visualization—Visualization techniques—; Human-centered computing—Visualization—Visualization design and evaluation methods

1 INTRODUCTION

The role of text in visualization is poorly defined. In most reference works on visualization the focus is on categoric, ordered or quantitative data as base types of data that are transformed into visual attributes then drawn in a display, e.g. [4, 12, 14, 31, 32, 56]. Text is usually understood as categoric data such as labels added on a graph or words in a tag cloud. However, categoric, ordered and quantitative encoding alone do not consider the unique role of literally encoding text. This paper will show functional, perceptual, cognitive, semantic and operational benefits that literal text affords visualizations.

Consider the portion of the world history chart in Fig. 1 [39]: it is a time-oriented visualization showing empires as streams: essentially a storyline visualization (e.g. [48] based on xkcd.com/657). However, it also has a large amount of text including the names of empires in big text, rulers, dates, key events, short descriptive phrases, full sentences and so on: the informational content far exceeds storyline visualizations and the encodings are beyond categoric data.

2 BACKGROUND

Literal text goes beyond categories or individual words.

2.1 More than categories

Categoric data in a visualization is typically transformed into a visual attribute such as hue or shape. However, perception limits the number of unique hues to ten or fewer [52]. While it is feasible to have an unlimited number of unique shapes [4], in practice very few unique shapes are provided by default in off-the-shelf software and programming libraries (e.g. nine in Excel, ten Tableau, seven D3.js).

There are tens of thousands of unique English words, with thousands understood by most adult speakers. There is a larger set of proper nouns. Then highly unique word sequences can be constructed to form phrases, sentences, stories, poetry and so forth. Further, words are much richer than categories and can express nuance e.g. many alternate words exist for each emotion [40].

*e-mail: rbrath@uncharted.software

Figure 1: Text-intensive chronological visualization of world history.

2.2 More than individual words

Combinations of words, relationships between words and sequences of words are typically lost in popular visualization approaches. Tag clouds typically represent individual words, e.g. [17]. Simple chart types, such as bar charts and pie charts, often use labels which require short text and tools often truncate text automatically. Some visualization techniques do work with longer text sequences such as WordTrees [54], Arc Diagrams [53] or some of the knowledge maps at scimaps.org (e.g. [8]) meaning that visualization of longer sequences do exist, but, these are not common.

Closely related shortened strings are the underlying analytic techniques that reduce text to individual words or phrases such as word frequency analysis, entity detection, sentiment analysis, emotion analysis, and vector-based representations such as Word2Vec (e.g. [34]). As such, the analysis of text by separating words loses much semantic context: Working with text in visualization means reconsidering text as a base type of data.

2.3 History

There are many historic visualization examples of text-based representations. Even within prose, highlighting techniques such as changes in color, italics, capitalization and weight have been used for centuries to make letters or words perceptually standout from surrounding text. And there are many other examples of text heavy visualizations such as:

Flowcharts document processes and may need text to describe each step, such as in Fig. 2 (from [10]).

Timelines may be heavily annotated such as Chapple and Garofalos *History of Rock N Roll* (in [51]); may contain names and phrases such as Figure 1; detailed sentences such as Rand McNallys *Histomap* [47], or many other examples in *Cartographies of Time* [43]).

Knowledge maps and *mind maps* require text, some of which contain large amounts of descriptive text such as Diderots *Tree of Knowledge* in Fig. 3 [44]. Modern procedurally-generated knowledge maps are text oriented as well, such as the proportional Euler diagrams in openknowledgemaps.org [28].

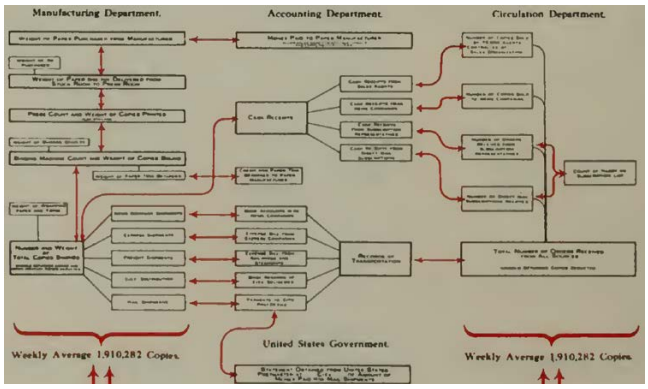

Figure 2: Portion of a text-intensive flowchart.

2.4 Information Visualization

In the field of information visualization, there are more than 400 text visualizations at the *Text Visualization Browser* (textvis.lnu.se). However, approximately one quarter of these visualizations have no text at all. Most of the remaining techniques are focused on labels, i.e. a word or two. Only a few manipulate individual characters (such as background per glyph [1] or font weight [36]). Some display longer text, although many of these limit visualization to manipulation of keywords in context (KWIC, e.g. [20]) and/or blocks of text; e.g. [46] does both.

Use of literal text, beyond simple labels, may be uncommon in information visualizations for many reasons [9], such as:

1. *Historic convention* separates words from images due to limitations of print-based processes or the extra expense associated with engravings, lithographs, etc.
2. *Display limitations* in early visualizations resulted in poor quality text (low resolution 96DPI) and limited real-estate (small screen sizes), which prompted guidelines to display details on demand [45].
3. *Graphic design conventions*, starting with mid-century modernism, favor less text, more white space and more graphics such as icons.
4. *Preattentive pattern perception* precludes reading text linearly.

These issues can be addressed, for example:

1. Modern web browsers do not have the same restrictions as print.
2. Modern displays are much higher resolution: a 4k display has 25x pixels of an early 1990's display.
3. Graphic design counter-trends bring more text manipulation into media, starting with post-modernists in the 1980's.
4. Some visualization tasks do not require preattention (e.g. inventories such as roadmaps), or subsets of text can be made preattentive with typographic formats [4, 5, 9].

Figure 3: Top part of Diderot's *Tree of Knowledge*.

Figure 4: Literal text benefits (red text) in relation to the visualization pipeline.

2.5 Characterization and Examples

The rest of this paper, and the key contribution, will characterize the benefits of literal text directly encoded in visualization and illustrate some examples focused on visualizations of strings of text.

3 Literal Encoding

Data can be encoded literally into strings of text and represented directly. Literal encodings are unique to text. Instead of considering text as a different form of data, some researchers consider text as a visual attribute, similar to shape, size or color, e.g. [7, 12, 32, 56], whereas other researchers do not include text. Bertin, for example, does not include text as a data type or encoding, however, Bertin narrowly defined his visual attributes as retinal variables, explicitly focusing on low level visual channels into which data is transformed for thematic representations for fast perception. Note that Bertin did discuss text and font attributes such as bold, italic and typeface in four pages in an appendix to the original French edition of Smiologie Graphique [4] and a followon article [5], which were not included in the later English translation.

There are many benefits to literal encoding, which can be organized by function, perception, cognition, semantics and operation. These benefits can be aligned to the *visualization pipeline* as shown in Fig. 4 (this diagram based on a simplified version from [12]).

3.1 Functional Benefits

Literal text may be effective in meeting the needs associated with the particular use case:

A. Information organization or communication purpose. The goal of a visualization is not necessarily analysis of patterns by preattentive perception, such as a pattern of dots in a scatterplot. Bertin defines other purposes, including communicating key information and organizing a large amount of data. Text is common in historic visualizations that organize large amounts of information as shown in the prior section. It is doubtful that these data-dense visualizations which organize information could function without text.

B. Text is the primary subject. Some visualizations are primarily about text. The simple tag cloud could not exist without text. Similarly, other visualizations related to text analysis presumably will have a strong need to explicitly represent literal text: extracted entities, topic analysis, sentiment and emotion analytics, keyword analysis, word stemming, dialogue analysis, taxonomies, and so on.

C. Increased information content. Many historic visualizations are highly similar to modern interactive visualizations. For example, Figure 1 is a storyline visualization. Flowcharts like Figure 2 are essentially graphs with much larger nodes sized to fit descriptive text. Fundamentally, visualization is a lossy medium [12], and the addition of text into visualization can increase the informational content.

D. High number of categories. Many visual attributes do not support categories of high cardinality as discussed in 2.1. Text can support high cardinality as words, phrases and sentences; or codes e.g. labeling aircraft in air traffic control, equipment in network diagrams, nodes in circuit diagrams, and so on.

E. Reduced design effort. It is feasible pictographic icons can be quick to decipher, but effective icons typically require significant effort by expert graphic designers to create a series of icons which work together.

F. Unambiguous Decoding. While an icon can represent a word or phrase, they can be ambiguous, e.g. Clarus the dogcow [18]. Text literally encodes the specific words, phrases or sentences of interest, so there is no information loss decoding back to text. Further, textual glyphs are highly learned and unambiguous in a well-designed font.

3.2 Perceptual Benefits: Fast Identification

Many visualization evaluations focus on testing time and errors, such as the rapid perception of the presence of a visual attribute. While text can be enhanced with attributes such as bold or italic to make a word preattentively pop-out, to read the text still requires active attention - which is slower than preattention.

However, real-world user tasks may be much more complex, requiring full decoding and comprehension of context. If the goal of a visualization is overall more efficient completion of tasks, then the performance of the entire system must be considered. Active attention of text can provide multiple perceptual benefits:

A. Reading is Automatic. Reading labels is very fast compared to interactions such as tooltips. The parallel letter recognition model of words states letter information is used to recognize the words [30], letters within a word are recognized simultaneously, and the and word perception is extremely fast [15].

Automatic word recognition is a common explanation for the Stroop effect which theorizes that reading is *automatic* and difficult to voluntarily stop. The *relative speed of processing model*, hypothesizes that it is faster to read the word than to name the color. Both explanations imply that reading words is fast. More generally, automaticity is the ability to perform a well practiced task with low attentional requirements: (a) the user is unaware that the task is occurring, (b) does not need to initiate the task, (c) does not have the ability to stop the process, and, (d) the task has low cognitive load [3]. The implication for visualization is that some text will be automatically read and that the cognitive cost of reading will be low.

B. Interaction is Slow. Interactions such as tooltips can be used to identify items on a screen. Interaction requires additional cognitive effort, as the viewer must first determine their interactive goal, engage in motor skills and progressively refine those skills to achieve the target [41].

Zoom and pan with level of detail for progressive appearance of labels is an alternative approach to revealing text, often used on interactive maps. While zoom-based labeling does occur in visualization (e.g. [26]), it is not common. Visualization does not have the equivalent labeling heuristics developed over centuries in cartography [42].

C. Legends are slow. In any form of visualization the viewer must both perceive the information of interest and then decode it. When visually assessing categoric data encoded as discrete instances in a visual attribute (e.g. using color to indicate categories), the viewer needs to recall the mapping between the colors and categories. With text, the viewer can directly decode the item.

D. Reduced load on short term memory. Tasks reliant on short-term human memory can benefit from explicit text. Text represented directly does not require short term memory, for example when referring back and forth between items and legends.

3.3 Cognitve Benefits: Aid to Reasoning

Beyond attention, problem-solving tasks require more cognitive resources to complete more complex tasks, such as determining the segment in a Venn diagram that corresponds to a logical condition, or, tracing a path in a network. Complex reasoning can be aided by the appropriate visual representations, as shown with force diagrams, pulley combinations and so forth (e.g. see many of the papers at *Diagrams* conference).

Readers construct mental representations of what they read at multiple levels: a) surface: words and syntax; b) propositional content; c) situation model incorporating and organizing content with respect to real world knowledge from memory [38].

A. Proposition and problem-solving tasks. The propositional model is self contained based on the presented facts: it can be understood through logical inferences, or through spatialization and visually queried. Problem solving is improved when textual instructions are directly integrated into diagrams [11]. Task performance improves when information is provided through both text and images with close spatial positioning and/or linkages [37]. More broadly, Larkin and Simons findings indicate that diagrammatic representations of content aid problem-solving [29] by:

1. *Locality:* Spatially organizing information together reduces search effort.
2. *Reducing cross-referencing:* Fewer steps to decode mappings.
3. *Perceptual enhancement:* Visual relationships support perceptual inferences.

Extending Larkin and Simons findings to text and visualization implies improved performance for text directly integrated into a visualization:

1. *Reduced Search (Locality).* Spatially organized displays aid finding information of interest. For example, in the knowledge graph (e.g. Fig. 3), the hierarchical relations and descriptions are spatially grouped facilitating search and providing text details. Identification of local relations is a common task in many visualizations and has been expressed by many researchers, for example:

   Tobler's first law of geography Everything is related to everything else, but near things are more related than distant things. [35] Many visualization layouts attempt to locate related objects close together (e.g. graph layouts, scatterplots, treemaps and set visualization). Text in any of these layouts supports identification and context of local related entities.

   Thundt et al's serendipity Serendipitous discovery is the fortuitous unexpected discovery by accident, often occurring during search. It is researched in library sciences and summarized by Thudt et al [49] who note that serendipity is closely associated with coincidence: wherein related ideas may manifest as simultaneous occurrences that seem acausal but meaningful. Both local proximity and related words and phrases can be recognized thereby providing different cognitive associations than visual associations alone.

2. *Reduced Cross-Referencing.* Text within a visualization element provides immediate detail as opposed to drill-down or referring to other linked visualizations:

   Tufte, Shneiderman and details Tufte popularized micro/macro readings; that is visual displays with high-density information, visually read at levels ranging from high-level macro patterns such as trends, clustering and outliers; down to localized patterns, such as individual observations and local peers [50]. Tufte summarizes this with to clarify, add detail. Tuftes approach is somewhat similar to Shneidermans visual information-seeking mantra: overview first, zoom and filter, then details on demand. [45] With Shneiderman, interactivity is used to reveal the detail information whereas Tufte plots it directly. Tufte plotted data on high-resolution paper, whereas Shneiderman was constrained to low-res 1990s screens, necessitating interaction.

   Bertin and categories Bertin states that shapes can represent categories of high cardinality and provides a map with 59 different glyph types (p. 157 in [4]). Bertin shows global patterns with respect to one shape cannot be seen. The viewer must linearly scan across glyphs and can only see local patterns. Furthermore, the viewer cannot ascertain if a local pattern is

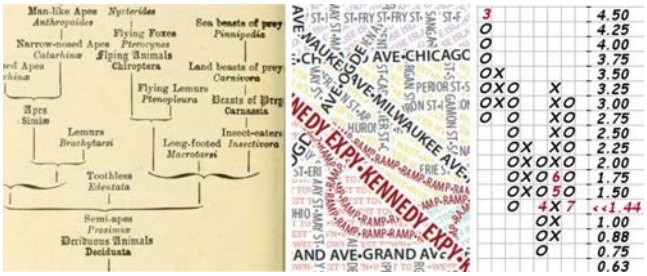

Figure 5: Perceptual inferences can be made via text density, lines and areas, or patterns forming columns.

meaningful without comparison to the global pattern. Text can be different. Text can be used non-categorically to encode unique entities and concepts (e.g. countries, people, idioms, quotes, etc.) a comparison to global is not needed. Furthermore, any unique text can be accessed non-linearly based on the layout, for example, to directly compare the unique text at two different locations of the plot area.

3. *Perceptual inferences* can be made across collections of labels or phrases. For example, a density of text can indicate common nodes across classification schemes in the graph by Haeckel (Fig. 5 left, [19]), the alignment of text along paths to form lines and areas (e.g. Axis Maps: Typographic Maps in Fig. 5 middle [2]), or the repetition of Xs and Os in point and figure charts (Fig. 5 right), e.g. [16].

3.4 Semantic Benefits: Word Sequences

Semantic content is added to text when words are combined in sequences. Consider the following:

- Jack and Jill went up the hill.
- Humpty Dumpty sat on a wall.
- The Owl and the Pussycat went to sea.

Text, when processed into discrete words, such as a word cloud, loses the semantic content of the word sequence. These sentences have meaning beyond the collection of words. Jack, Jill and Humpty all have height and might fall. Jack, Jill, the Owl and the Cat went somewhere, Humpty did not.

Analytic and visualization approaches need to consider applications where word sequence is maintained, or assemblies of related words. Search user interfaces have evolved significantly beyond simple keywords and document metadata to include contextual titles, phrases and sentences in search results (e.g. [20] Chapter 5). Studies have shown superior performance for results containing text snippets [13].

More broadly, a variety of visualization techniques have evolved in the humanities for close reading, summarized by Jnicke et al [25]. For close reading, these approaches maintain word sequence and change visual attributes to markup the text (similar to keyword markup), superimpose markers and connections on top of words (e.g. like a graph), or increase space between letters and words to add markers such as flows or phonetics (e.g. like WordTree).

3.5 Operational Benefits: Text Manipulation

Literal text allows for textual operations which can be used to preprocess text, or interactively to operate on text to provide for additional uses:

- *Ordering.* Text can be ordered and sorted (alphabetically), facilitating search and lookup.
- *Search and Filter.* Text can be navigated and/or filtered by complete or partial strings.
- *Summarization.* Text can be summarized: longer texts can be reduced to shorter texts while retaining a subset of the original meaning.
- *Comparison, similarity and translation.* Text can be compared and assessments made regarding similarity. Tools such as thesauruses and dictionaries aid in understanding similarity. Translation extends comparative analysis to produce near equivalent meaning in another language.
- *Tone, opinion, sentiment and emotion.* Text can be evaluated assigned quantitative values, for example, for opinion (e.g. ratings), tone (e.g. news tone), sentiment (e.g. score for positive or negative), or emotion (e.g. based on text or other cues).
- *Categorization, taxonomies and topic analysis.* Organization of many texts follow classification schemes and tagging of topics by keywords; such as general purpose classifications (e.g. Dewey, Library of Congress) and domain specific classifications (e.g. ACM, IEEE).
- *Natural Language Processing and Machine Learning.* All of the above operations are rapidly evolving with NLP and machine learning, such as topic extraction, emotion analysis, sentiment detection, machine translation, automated summarization and so on. For example, Zhang et al.'s recent transformer models have achieved near human-level abstractive summarization [57]. Open source NLP libraries, such as NLTK, spaCy, and compromise [6, 23, 27], provide parts of speech tagging, tokenization, entity extraction, dependency parsing, and entity linking, which can be assembled into a wide variety of textual analyses extracting and linking related words and phrases, such as such as character descriptions or rhetorical devices.

4 Discussion

There are many ways to consider these benefits in the humanities. As there are many existing visualization techniques focused on collections of disconnected words (e.g. word clouds), the following focuses on applications relevant to semantic word sequences (i.e. phrases and sentences) and the corresponding benefits.

A. Sentences on Paths. In Fig. 6, Nigel Holmes reduced key testimony in the Iran-Contra affair into a set of statements implicating individuals represented as color-coded textual connections between politicians (i.e. a graph) [22]. Literal text is required for the communications purpose (3.1:A). Magazines are non-interactive the text is a more efficient alternative to referring to a legend (3.2:C), and aids reasoning with directed connections enabling perceptual inferences (3.3:A3).

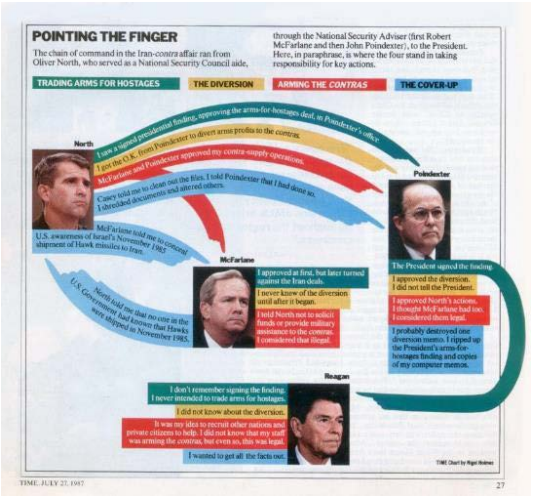

Figure 6: Connections between individuals with testimony.

Figure 7: Portion of Amsterdam SmellMap with narrative paths.

Kate McLean uses linear text along geographic contour lines to form artistic geospatial narratives in Fig. 7 [33]. The literal text is the primary subject (3.1:B), which the viewer will automatically start to read (3.2:A), and compare local information within a geolocation (3.3:A1), with added potential for residents to supplement real-world local knowledge (3.3:B).

Programmatic examples include news headlines with associated stock price movement, and twitter commentary with associated retweet counts [9]; which enable comparison between content (the text) and the response (the behaviour of markets / sharing by users).

B. Small Blocks of Text. *NewsMap* extends the visualization layout of a treemap with news headlines Fig. 8 [55]. Information content beyond a treemap is increased with hue, brightness and text (3.1:C). Interaction is not a suitable: direct reading is required (3.2:B). The treemap layout of text ensures that similar headlines are spatially proximate (3.3:A1) and augments the viewers real-world understanding of news (3.3:B). Interactive search and filter aid exploration (3.5).

Figure 8: NewsMap headlines. Size, hue and brightness encode data.

One challenge with the treemap layout is the prioritization of area accuracy over the display of text, as such, some squares are too small for legibile text, or text does not fit well. Consider a dataset of 2894 coroners inquests from Georgian London [21, 24] where each case is summarized by a sentence and verdict, e.g.

- Mary Roberts drowned herself. Suicide.
- Mary Gardiner struck with hand. Homicide.
- Ann Fitsall suffocated and burnt. Accident.
- Nicholas Bone, John Dayson and James Cusack killed by a brick wall. Accident.

An analysis may be interested in causes of death. From each case can be extracted a subject, verb and object. These can be organized into a hierarchy, by verb (e.g. drowning, suffocation), object (e.g. hand, wall) and subject (e.g. Mary Roberts, Ann Fitsall). There are a few caveats, for example, where two or more verbs or objects are used in a single sentence, only one is selected so that individuals are not double counted. The hierarchy can be visualized, for example, with a treemap (Fig. 9), which draws attention to big boxes and bright colors (e.g. drowned) but at the cost of fitting labels or skipping labels on small boxes.

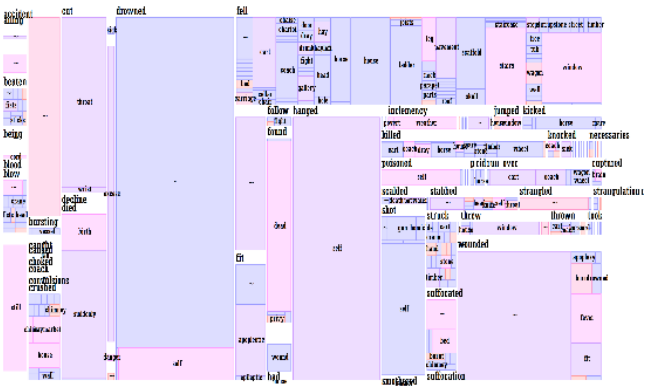

Figure 9: Treemap showing causes of death in Georgian London.

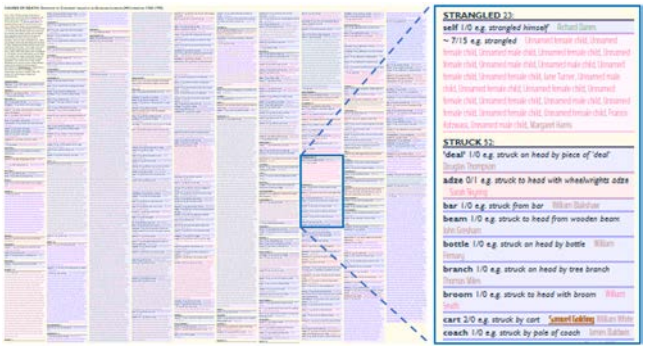

Figure 10: Activities, objects, and the names of deceased.

Instead, Fig. 10 shows the same data in a representation conceptually similar to a dictionary listing with font weight, color and background shading conveying data, live at `https://codepen.io/Rbrath/full/rNewgde`. All text is visible and readable on a 4K display: an enlarged portion is shown at right (i.e. organizational purpose 3.1:A). Broad patterns are visible, for example, large blocks of names such as the red text (homicide) on red background (gender female) under the verb **STRANGLED** (i.e. a preattentive perception of the large block, then the bold allcaps text, which is identified by automatic reading 3.2:A). On close inspection, the local text can be read, for example, many local objects associated with the verb **STRUCK** (e.g. **adze, bar, beam, bottle, broom**) and the names of the victims (e.g. **adze**: Sarah Skyring, **bar**: William Blakshaw, etc) no drill-down is needed (3.3:A2). Furthermore, search can be used to find the common forms of death for child, or filters can be used to isolate only homicides (3.5).

Fig. 11, is a simple adjacency matrix indicating dialogue from one character (vertical axis) to another character (horizontal axis), with bubble size proportional to the amount of dialogue.

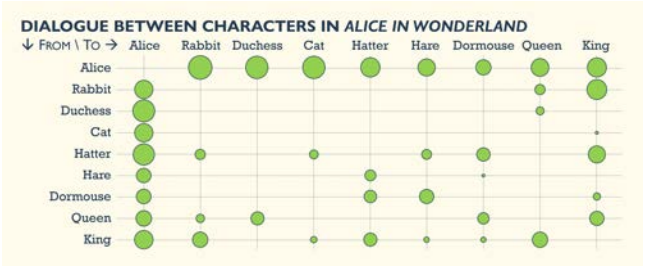

Figure 11: Volume of character dialogue via an adjacency matrix.

Figure 12: Character dialogue via an adjacency matrix, with highlighted repeated sets of words per character.

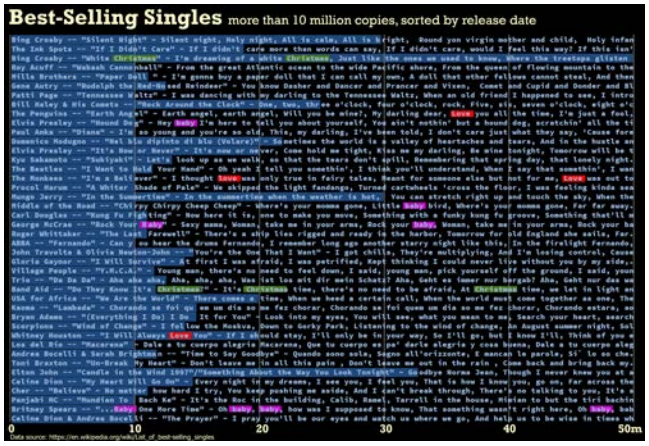

Figure 13: Song title and lyrics over bar chart and repeated words.

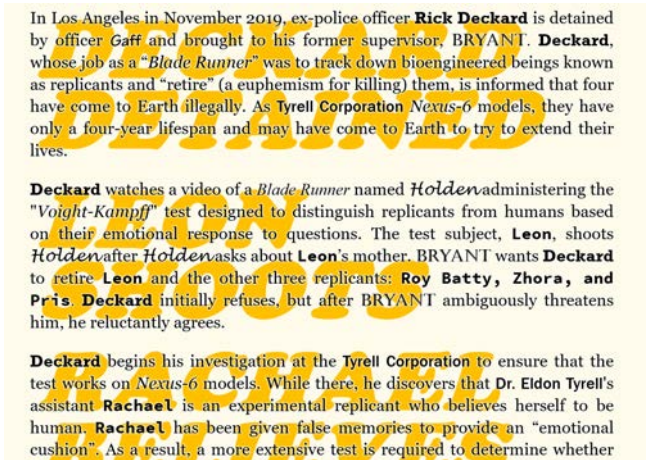

Figure 14: Body text over a two-word paragraph summary.

Fig. 12 is the same adjacency matrix where each cell contains the dialogue from one character to another, increasing content (3.1:C). Cells are expanded if the next cell to the right would otherwise be empty. The matrix contains much of the dialogue, although some cells are truncated. In addition to showing dialogue, some phrases have been highlighted: an simple NLTK algorithm has tagged sets of words that are frequently repeated by characters (not including sets made of only prepositions, articles and conjunctions), so for example, the moral of that is is highlighted each time the Duchess repeats the phrase to Alice (aiding perceptual inferences 3.3:C). Word sets are repeated words, but the order may vary. This also highlights Carrolls logical inversions, such as the Hatters I see what I eat and I eat what I see. It also highlights variants of catchphrases, such as off with his/her head yelled by the Queen. Highlighting similar words in different orders aids analysis of the variation of the semantic content (3.4). This use of NLP could be further tuned, for example, to tag rhetorical devices; which in turn, could be visualized in techniques such as an adjacency matrix, story flow, text on path, and so on.

C. Superimposition. Text can be superimposed over top of data-driven graphical marks similar to using a highlighter, where the highlight may represent quantities or keywords or so on.

Each row in Fig. 13 is a top selling song, shown as a line of text identifying artist, song tile, and opening lyrics. The line is superimposed over a blue bar chart indicating number of singles sold and also superimposed over highlights of common keywords such as Christmas, love, and baby. The superimposed text per bar increases information beyond the standard bar chart (3.1:C); provides detail instead of an interactive tooltip (3.2:B), allows perceptual inferences (patterns of repeated words) (3.3:A3) and reading can trigger real world recall of a familiar tune (3.3:B).

Fig. 14 shows body text superimposed over summary words extracted via NLP. When reading a long text on a mobile device only a paragraph or two are visible, meaning that traditional textual cues such as headings and spacing may be off screen, thereby making it more difficult to navigate around the text. Using a large proper noun and associated verb could facilitate skimming while scrolling by providing textual landmarks to characters and their actions. This large text could disappear on scroll-stop. This is a summary operation (3.5) to reduce text, arguably to a high number of categories (3.1:D); and reducing cognitive load by not requiring the viewer to recall off-screen headings (3.2:D).

5 Conclusion

The itemized literal benefits provide a starting point for researchers and designers to consider the alternative visualizations and the potential benefits across different design alternatives. These literal benefits can be considered in any kind of text-enhanced visualization, whether simple labels, or through to the examples shown here.

The example analyses here focus on collections of short semantic word sequences. The examples here may be highly-specific, however, the intent is to illustrate the possibilities and benefits of literal depictions. These should encourage digital humanities researchers to search for new visualization techniques. It should also encourage visualization researchers to validate benefits beyond prior research in adjacent fields through experimentation directly on text-orientated visualizations to quantify and characterize benefits. And, it should encourage linguistics and NLP researchers to innovate with humanities researchers to find and annotate word sequences, which can be used in human-in-the-loop analytical user interfaces for use by digital humanities scholars.

Acknowledgments

Images created by the author are CC-BY-4.0. Some of these images and portions of the text will appear in the forthcoming book *Visualizing with Text*, by Richard Brath, as part of the Visualization Series edited by Tamara Munzner to be published by A.K. Peters in November 2020.

References

[1] A. Abdul-Rahman, J. Lein, K. Coles, E. Maguire, M. Meyer, M. Wynne, C. R. Johnson, A. Trefethen, and M. Chen. Rule-based visual mappings–

with a case study on poetry visualization. In *Computer Graphics Forum*, vol. 32, pp. 381–390. Wiley Online Library, 2013.

[2] S. Afzal, R. Maciejewski, Y. Jang, N. Elmqvist, and D. S. Ebert. Spatial text visualization using automatic typographic maps. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12):2556–2564, 2012.

[3] J. A. Bargh. The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In *Handbook of Social Cognition*, pp. 1–40, 1994.

[4] J. Bertin. *Sémiologie Graphique*. Gauthier-Villars, Paris, 1967.

[5] J. Bertin. Classification typographique: Voulez-vous jouer avec mon a. *Communication & Langages*, 45(1):70–75, 1980.

[6] S. Bird, E. Klein, and E. Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009. See also: `www.nltk.org`.

[7] K. Börner. *Atlas of Knowledge: Anyone Can Map*. MIT Press, Cambridge, Mass, 2015.

[8] K. Börner, E. Hardy, B. Herr, T. Holloway, and W. B. Paley. Taxonomy visualization in support of the semi-automatic validation and optimization of organizational schemas. *Journal of Informetrics*, 1(3):214–225, 2007.

[9] R. Brath. *Visualizing with Text*. A.K. Peters, Boca Raton, FL, 2020.

[10] W. C. Brinton. *Graphic Presentation*. Brinton Associates, 1939.

[11] P. Chandler and J. Sweller. Cognitive load theory and the format of instruction. *Cognition and instruction*, 8(4):293–332, 1991.

[12] M. Chen and L. Floridi. An analysis of information in visualization. *Synthese*, 2013.

[13] C. L. Clarke, E. Agichtein, S. Dumais, and R. W. White. The influence of caption features on clickthrough patterns in web search. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 135–142. ACM, 2007.

[14] W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554, 1984.

[15] M. Davis. Aoccdrnig to a rscheearch at cmabrigde... `http://mrc-cbu.cam.ac.uk/people/matt.davis/Cmabrigde`, 2003.

[16] V. DeVilliers and O. Taylor. *The Point and Figure Method of Anticipating Stock Price Movements*. Harriman House, 1933.

[17] J. Feinberg. Wordle. *Beautiful Visualization: Looking at Data Through the Eyes of Experts*, 2010.

[18] S. Hackett. The history of clarus the dogcow. `http://512pixels.net/dogcow/`, 2015.

[19] E. Haeckel. *The Evolution of Man: A Popular Exposition of the Principal Points of Human Ontogeny and Phylogeny*. D. Appleton and Company, 1897.

[20] M. Hearst. *Search User Interfaces*. Cambridge University Press, 2009.

[21] T. Hitchcock, R. Shoemaker, S. Howard, J. McLaughlin, and et al. London lives, 16901800. `www.londonlives.org`, 2012.

[22] N. Holmes. Pointing the finger. In *Time Magazine*, p. 27, 1987.

[23] M. Honnibal and M. Johnson. An improved non-monotonic transition system for dependency parsing. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pp. 1373–1378, 2015. See also: `spacy.io`.

[24] S. Howard. London lives coroners inquests, 17601799. `london.sharonhoward.org/CoronersInquests`, 2018.

[25] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. On close and distant reading in digital humanities: A survey and future challenges. In *Eurographics Conference on Visualization (EuroVis)-STARs*, pp. 83–103, 2015.

[26] D. Jonker, S. Langevin, D. Giesbrecht, M. Crouch, and N. Kronenfeld. Graph mapping: Multi-scale community visualization of massive graph data. *Information Visualization*, 16(3):190–204, 2017.

[27] S. Kelly. Compromise: modest natural language processing. `https://github.com/spencermountain/compromise`, 2019.

[28] P. Kraker, C. Kittel, and A. Enkhbayar. Open knowledge maps: Creating a visual interface to the worlds scientific knowledge based on natural language processing. *Zeitschrift fr Bibliothekskultur/ Journal for Library Culture*, 4(2):98–103, 2016.

[29] J. H. Larkin and H. Simon. Why a diagram is (sometimes) worth ten thousand words. *Cognitive science*, 11(1):65–100, 1987.

[30] K. Larson. The science of word recognition. *Advanced Reading Technology, Microsoft Corporation*, 2004.

[31] J. MacKinlay. Automating the design of graphical representations. *ACM Transactions on Graphics*, 5, 2006.

[32] R. Mazza. *Introduction to Information Visualization*. Springer-Verlag, 2009.

[33] K. McLean. *Smellmap Amsterdam*. Sensory Maps, 2014.

[34] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint*, 2013.

[35] H. Miller. Toblers first law and spatial analysis. *Annals of the Association of American Geographers*, 94:284–289, 2004.

[36] M. Nacenta, U. Hinrichs, and S. Carpendale. Fatfonts: combining the symbolic and visual aspects of numbers. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, vol. 57, pp. 407–414, 2012.

[37] N. H. Narayanan and M. Hegarty. Multimedia design for communication of dynamic information. In *International journal of human computer studies*, vol. 57, pp. 279–315, 2002.

[38] S. Payne. Mental models in human-computer interaction. In *The Human-Computer Interaction Handbook*, pp. 80–93. CRC Press, 2007.

[39] E. Pick. Tableau de lhistoire universelle. (eastern hemisphere). `davidrumsey.com/luna/servlet/s/61kab2`, 1858.

[40] R. Plutchik. The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American Scientist*, 89:344–350, 2001.

[41] R. W. Proctor and K.-P. L. Vu. Human information processing: an overview for human-computer interaction. *The human-computer interaction handbook*, 89:60–79, 2007.

[42] A. Robinson, J. Morrison, P. Muehrcke, A. Kimerling, and S. Guptill. *Elements of Cartography*. John Wiley & Sons, New York, NY, 1995.

[43] D. Rosenberg and A. Grafton. *Cartographies of Time: A History of the Timeline.* Princeton Architectural Press, 2012.

[44] C. Roth. Genealogical distribution of the arts and sciences. `omeka.lehigh.edu/exhibits/show/diderot/sciences/tree`, 1780.

[45] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE symposium on visual languages*, pp. 336–343. IEEE, 1996.

[46] M. Skeppstedt, K. Kucher, M. Stede, and A. Kerren. Topics2themes: Computer-assisted argument extraction by visual analysis of important topics. In *Proceedings of the 3rd Workshop on Visualization as Added Value in the Development, Use and Evaluation of Language Resources at LREC*, vol. 18, pp. 9–16, 2018.

[47] J. Sparks. The histomap. four thousand years of world history. relative power of contemporary states, nations and empires. `davidrumsey.com/luna/servlet/s/61grv9`, 1931.

[48] Y. Tanahashi and K.-L. Ma. Design considerations for optimizing storyline visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2679–2688, 2012.

[49] A. Thudt, U. Hinrichs, and S. Carpendale. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1461–1470. ACM, 2012.

[50] E. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1983.

[51] E. Tufte. *Envisioning Information*. Graphics Press, 1990.

[52] C. Ware. *Information Visualization: Perception for Design*. Springer-Verlag, 2000.

[53] M. Wattenberg. Arc diagrams: Visualizing structure in strings. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002.*, pp. 110–116. IEEE, 2002.

[54] M. Wattenberg and F. B. Viégas. The word tree, an interactive visual concordance. *IEEE transactions on visualization and computer graphics*, 14(6):1221–1228, 2008.

[55] M. Weskamp. Projects: Newsmap. `marumushi.com/projects/newsmap`, 2004. accessed: 03/03/2014.

[56] L. Wilkinson. *The Grammar of Graphics*. Springer, 1999.

[57] J. Zhang, Y. Zhao, M. Saleh, and P. J. Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *Thirty-seventh International Conference on Machine Learning*, 2020.