

Visualization of Narrative Structure

Natalia Y. Bilenko

Helen Wills Neuroscience Institute

University of California, Berkeley

Berkeley, CA 94720

Email: nbilenko@berkeley.edu

Asako Mivakawa

Helen Wills Neuroscience Institute

University of California, Berkeley

Berkeley, CA 94720

Email: amiyakawa@berkeley.edu

Abstract—In this work, we visualized two aspects of narrative structure: character interactions and relative emotional trajectory. We performed the analysis for three different books: *The Glass Menagerie* by Tennessee Williams, *Kafka on the Shore* by Haruki Murakami, and *The Hobbit* by J.R.R. Tolkien. We visualize character relationships in a dynamic graph that displays the evolution of connections between characters throughout the book. We show emotional polarity and intensity in a color-coded sentiment bar plot that provides access to the original text. The emotional path of each character through the book can be traced by selecting a character, which highlights the sentences in the sentiment plot where this character appears. Our result is a visualization that provides an at-a-glance impression of the book as well as a method for exploring the plot details. We hope that this method of visual book summarization will be useful both for commercial applications such as book shopping or e-readers, as well as for succinct literary analysis.

INTRODUCTION

According to the Association of American Publishers 2012 report, digitized literature (e-book) sales grew 45% from 2011 and constitute 20% of trade revenue in 2012. Self-published e-book is one of the fastest growing sectors in the publication market (source: Bowker Market Research). With increasing volume and availability of digitized literature, there is enormous demand for automated text analysis and visualization. Analysis of narrative structure is one area of textual analysis that has proven difficult to automate and display, due to high dimensionality and challenges of natural text processing, as well as the abstraction of visual data and difficulty of mapping text to visual elements. Parsing and tagging of text-based data can be tedious, difficult, and time-consuming. Additionally, currently available text visualization techniques are limited. For example, Word Clouds, a popular method of text visualization, eliminates contexts and use non-quantitative representations (see Figure 1). Since text itself is visual, summarizing and reducing it can be quite difficult. Often, text-based representations are visually crowded and overwhelming (see Figure 2).

Nonetheless, visualizing the narrative of a book is useful for commercial and literary analysis purposes. On one hand, visualization of book contents could provide a quick overview of the book. For example, displaying the emotional trajectory of a book as a progression of color patches could help a reader quickly make an informed book selection on a shopping website. It could also be useful as an e-reader application, to help the user keep track of their progress in a book. On the

other hand, visualization can provide a novel perspective into the structure of the book as a form of brief literary analysis. We aimed to represent narrative of a fictional book via character interactions and emotional trajectory. Though those features are far from a comprehensive summary of narrative, they shed a light on the progression of a story. The live demo of the visualization is available at <http://tiny.cc/narrative-viz>, and the source code can be found at <https://github.com/nbilenko/narrative-viz>.



Fig. 1: Visualization of the Obama Nomination Acceptance Speech produced using Wordle (<http://www.wordle.net>)




Fig. 2: Text Arc visualization of *Alice’s Adventures in Wonderland* by Paley. (1991)

RELATED WORK

Sentiment Analysis

Previous textual analysis work has focused on analysis of relatively short texts such as consumer reviews (Hu and Liu, 2004, Socher et al., 2013), twitter feeds (Twitter StreamGraphs) and emails (Cheng et al., 2011). The goal of these analyses was correct automated identification of polarizing features, such as user ratings of a movie, sentiment

of a twitter feed, or gender identity of e-mail authors. However, compared to the short texts used in these analysis, narratives of fictional books pose additional challenges. These narratives tend to be less polarizing on a shorter timescale, and sentiments in them tend to evolve slowly and more subtly. Additionally, these narratives are usually written from multiple perspectives of different characters, rather than a singular perspective. Therefore, narrative speaker needs to be known to correctly interpret the sentiment analysis results.

Visualization of fictional narratives

Some researchers have worked specifically on visualizing fictional narratives. These works include Stefanie Posavec's *On the Road* (see Figure 3), W. Bradford Paley's *Alice's Adventures in Wonderland* (see Figure 2), Matthew Hurst's *Pride and Prejudice* (see Figure 4). Each of these work highlighted new aspects of the structure of these books. Posavec and Hurst created their visualizations as static symbolic chart with graphic features (e.g. color, length, size, shape) representing literary features (e.g. theme, character occurrence, or sentence length). Their visualizations highlighted coarse structure of a book and may be useful for comparisons across different books. However, their visualization eliminated all of the original texts and gives little textual contexts to the visualization. Posavec's visualization in particular was constructed by hand and thus makes further investigation difficult. Paley took a finer-grain approach and used parts of raw text in his visualization. Word co-occurrences are represented as line connections between words and text highlights that dynamically change upon user interaction. While users may be able to extract meaningful literary features like character relations or recurring themes, information density of this visualization is quite high and some visual parameters (e.g. overwrapping texts in small fonts) are potentially overwhelming.

Openbible.info created a static visualization focusing on a sentiment analysis of the Bible (see Figure 5). In order to provide overall sentiment, rather than sentence-sentence sentiment assignment, the authors took a moving average of either 5 or 15 verses and presented the average results chronologically. Some chronological contexts are provided as a radial segments representing main events (e.g. exile), the book and chapter ID (Luke, John etc.) and the first verse of the chapter.

There has been some work in representing interactions within stories. For example, Sudhahar et al. (2011) use Quantitative Narrative Analysis to find relationships between actors in a corpus of news stories. They further characterize these relationships by identifying the verbs that categorize the interactions (see Figure) They find that their results reflect the representative stories that appeared in the news cycle at that time. This type of character interaction analysis can be relevant for fiction visualization as well. However, classifying the actions between characters is often more difficult in fiction. The verbs that describe character interactions in a fictional book are often less informative and usually relate to dialogue (such as "said", "told", etc.). On the other




Fig. 3: Manual visualization of *On the Road* by Stefanie Posavec, "Writing Without Words"




Fig. 4: *Pride and Prejudice* visualization by Matthew Hurst, "Visualizing Jane Austen" (2011)

hand, in news stories, action verbs are usually used to maximize informational content, and dialogue replication is very uncommon. Due to these challenges, co-occurrence of characters in fiction is often a more informative metric of interaction than action analysis.

I. METHODS

We propose a method for visualizing the narrative of a fictional book in terms of its emotional trajectory, character relationships, and overall structure. We use a combination of natural language processing methods and d3.js (Data-Driven Documents) visualization tools to capture different aspects of the narrative.

A. Texts

We used three texts for our analyses: *The Glass Menagerie* by Tennessee Williams, *Kafka on the Shore* by Haruki Murakami,




Fig. 5: Visualization of the Bible, openbible.info

Fig. 6: News story interactions via Quantitative Narrative Analysis

and *The Hobbit* by J.R.R. Tolkien. We chose these three books because they have variable length and number of characters, as well as the structure of the story, intended audience, and language complexity. *The Glass Menagerie* has 2478 sentences, 7 acts, and only 4 characters, and it is a haunting memory play. *Kafka on the Shore* has 16413 sentences and 14 major characters, and it is a metaphysical magical realism novel. Finally, *The Hobbit* has 5892 sentences and 36 major characters, and is a children’s fantasy novel.

B. Sentiment Analysis

To capture the emotional trajectory as it develops throughout the book, we use Treebank Sentiment Analysis to quantify emotion dynamically throughout the book on a paragraph scale. Treebank Sentiment Analysis is part of the Stanford Natural Language Processing Groups processing software. Using the Stanford Parser, input texts are first broken up into words and sentences (tokenization), different forms of words are grouped together (lemmatization) and sequential tokens are clustered into sentences (sentence splitting). Grammatical relations of tokens in each sentence are analyzed and annotated accordingly (parsing). Sentiment Treebank algorithm then creates a hierarchical tree that represents compositional structure of a sentence and assigns sentiment label to each tree node. Since this model is sensitive to compositional relations between words, it outperforms traditional bag-of-words model, which consider each word in isolation. Socher et al. estimate the overall classification accuracy of a single sentence to be 85.4%, which is 9.7% above the bag-of-words model. The algorithm learns in a supervised manner by manual labeling of phrases by humans. We used the version of algorithm that is already trained on 215,154 unique phrases from 11,855 single sentences (derived from rottentomatoes.com movie reviews), annotated by 3 human judges. Using the output of Sentiment Treebank, we will visualize the emotional trajectory by encoding sentiment




Fig. 7: Moviebarcode visualization of *Finding Nemo*, *Amelie*, and *Scarface* (<http://moviebarcode.tumblr.com>)

as color and displaying the change in sentiment on the chronological axis of the book. We plan to use a visualization technique similar to the Movie Barcode (see Figure 7 for the Fashion Fingerprints 8, except the color would represent the emotional content of the book instead of the actual color of movie frames or photographs.

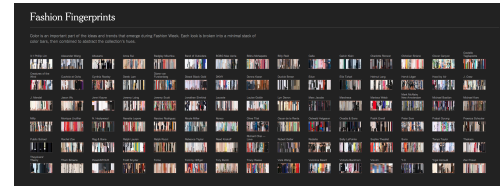


Fig. 8: Fashion fingerprints visualization of New York Fashion Week in New York Times by Bostock et al. (2013)

C. Character interactions

We visualize the evolution of relationships between characters in the book. We use the Natural Language Toolkit (nltk.org) to extract character appearances and co-occurrences in each chapter of the text. Co-occurrences are defined as occasions when two characters appear within a specified neighborhood (11 sentences) within each chapter. We then visualize the characters as nodes in a graph with co-occurrences between them as edges. The thickness of the connections represents the number of co-occurrences between characters in each chapter. The graph changes dynamically to represent changing character interactions as the user scrolls through chapters of the book.

D. Combined visualization

To combine the two aspects of visualization, we make it possible to highlight only sentiment-rated sentences where a single character appear by selecting that character’s name. In the complete visualization, the user sees an interactive chronological axis of the book with chapters marked. The emotional progression displayed as a color bar graph is a second chronological axis. The character relationship graph represents the character interactions for the currently selected chapter in the book.

RESULTS

Our visualization started with an introduction screen that gave users a brief purpose and overview of the visualization. Users were prompted to select one book to visualize by clicking a button. Three sample texts used in our visualization were *The Glass Menagerie*, *The Hobbit*, and *Kafka on the Shore*. After selecting the book and upon entering the visualizer,

the user saw a pop-up instruction box that described three main features on the screen: the chapter slider, the character interaction graph, and the sentiment bar plot (see Figure 9). Large screenshots of all three visualizations are at the end of the paper (see Figures 10, 11, 12).

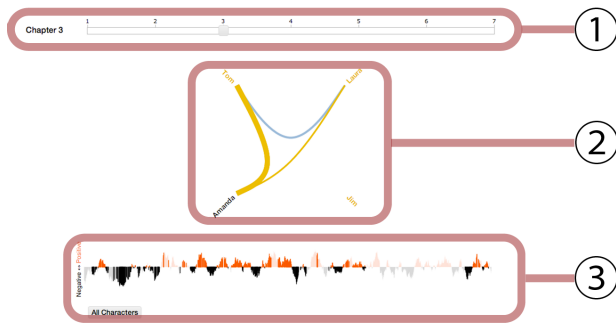


Fig. 9: Visualization layout with (1) the chapter slider, (2) the character interaction graph, and (3) the sentiment bar plot.

Feature 1: Character interaction graph with slider

Character interaction graph showed all the characters in the book in a radial layout. The links between the characters showed the co-occurrences of those characters in the selected chapter. Characters were grouped roughly according to the character type if such groupings were applicable for the story. For example, in *The Hobbit*, there were several character types: hobbits, elves, dwarves, et cetera. Character groupings were used to create another layer in the hierarchical edge bundling, which forces each edge to take a different arc path depending on whether the nodes its connecting are in the same or different group. For instance, an edge connecting two hobbits takes a shorter and flatter arc than one connecting a hobbit and an elf, which will travel through the center. This feature allowed users to easily observe whether the interactions that took place in a chapter was within or across character types. The thickness of edges represented the number of co-occurrences between two characters in each chapter. Users moved the chapter slider to observe the chapter to chapter changes of character interactions. We observed interesting features, like graphs changing from intragroup to intergroup connections (see Figure 13) or alternating pattern of character connections (see Figure 14).




Fig. 13: Evolution of character interactions in *The Hobbit*. Note that the switch from intragroup to intergroup connections.

Feature 2: Sentiment bar plot

The sentiment bar plot is located at the bottom of the page. Each bar represent a sentence in a novel organized in a




Fig. 14: Evolution of character interactions in *Kafka on the Shore*. Note that the chapters alternate between two sets of characters.

chronological order. Through the sentiment bar plot, user could access the original sentence that gave rise to the sentiment rating by hovering over the bar (see Figure 15). The sentiment bar is roughly aligned with the chapter slider at the top of the page, however, we did not correct for the different chapter lengths in the current version. The height of the sentiment bar shows the sentiment polarity (positive vs. negative) and the emotional intensity as a distance from zero, defined as the average sentiment polarity of the book. Plots of raw sentiment values at single sentence level were dominated by large local value fluctuations and interpretation was difficult. To correct this issue, we calculated the moving average with set bins (5 sentences on each side for *Kafka on the Shore* and 20 sentences on each side for *The Hobbit* and *The Glass Menagerie*). This smoothing procedure resulted in smooth peaks in both directions that are designed to corresponds roughly with scenes in the case of *The Glass Menagerie*, and paragraphs in the case of *The Hobbit* and *Kafka on the Shore*. Current smoothing parameters could potentially be adjusted. Other methods such as taking average within an actual paragraph or specific topic contexts for instance might affect the accuracy of sentiment representation and could be tested (see Discussion).

Sample positive sentence:



Sample negative sentence:




Fig. 15: The users can access the original sentence that gave rise to the sentiment rating by hovering over the bar.

Feature 3: Exploring character contributions to a scene sentiment

Users could highlight the sentiment bars associated with a certain character by clicking the character name in the interaction graph. Upon clicking the name, the sentiment bars

associated with a chosen character increased in opacity and the bar width, resulting in a higher visibility. This feature could be used for revealing the overall sentiment associated with the character. For example, in *The Glass Menagerie*, the sentiment of the sentences where Laura appears is skewed toward the negative polarity, while sentences in which Jim appears tend to be positive. Using the character interaction graph, users are able to first identify the chapter the story where two characters interact, and then compare their sentiments by using character highlight features in the sentiment bar plot (see Figure 16). This type of sentiment comparison may provide additional insight into character relations.




Fig. 16: The users can first identify the chapter the story where two characters interact, and then compare their sentiments by using character highlight features in the sentiment bar plot.

User Experience

We observed the experience of the users during the CS294 final project poster session. Users were given an in-person tutorial by one of the authors and prompted to test out the program. User seemed to navigate well with basic common features, like toggling chapter slider, without an explicit instruction. Therefore, our feature 1 seems to be most intuitive of all features. Less common features, like mouseover to reveal the original sentence, seemed to require an extra verbal prompt by the experimenter. However, once familiar with these features, the users spent a significant amount of time clicking on the character to highlight the bar and hovering over to reveal the sentences. Many users commented positively on the accessibility to the original text. Several users liked the ability to select the sentences in the sentiment bar plot that related to a single character, commenting that they appreciated the integration. In general, users preferred to see the visualization with the book they have read, and tended to stay in one visualization, rather than switching between books.

DISCUSSION

Our sentiment analysis was done on a single sentence level and later averaged across neighboring sentences. Although this approach allows easier interpretation of the overall sentiment in the scene, it is problematic that this method might average out a sentiment that happens on a shorter time scale (a few sentences). More appropriate approach is to segment narratives into literally meaningful units such as paragraphs, chapters, or topics and take a running average within these units. This could be achieved by extracting meaningful labels from the text, like the word chapter, or using hand-segmented text. Alternatively, we could try automated segmentation methods,

like Bayesian unsupervised topic segmentation method proposed by Eisenstein and Barzilay (2008). Furthermore, we have noticed that sentiments could vary significantly between different characters appearing in the same scene. Labeling each dialogue with the character and taking sliding average within character may also increase the accuracy of the sentiment analysis. One benefit of the Sentiment Tree Bank is that it uses machine learning techniques to update sentiment representations. We used the algorithm that was pre-trained on Rotten Tomatoes movie review sentiment classification. It is possible that some of the sentiment was misclassified due to the stylistic differences between fiction novels and the training sets. It will be interesting to test if the algorithm will perform better if trained on sentences that has similar structures as the genres of the novel or the specific author.

FUTURE WORK

Our method for visualization of narrative can be expanded in many potential directions. First, our definition of character relationships was based on simple co-occurrence. Using a more sophisticated analysis, such as action-based grammatical analysis of character relationships could provide more context for interactions. Additionally, a topic analysis could provide a better map of the situations in which characters interact. The correspondence between chapter boundaries and the sentiment axis of the story is loose in our visualization. Labeling chapter boundaries and aligning them with the sentiment bar plot would improve the analysis of this aspect of narrative. Exploring the relationship between sentiment fluctuations and event boundaries in the story is an important extension of our research.

We used a fairly simple binary sentence-based sentiment analysis in this work. Incorporating context and using larger neighborhoods than single sentences to evaluate sentiment would likely improve the accuracy and precision of our ratings, as well as limiting the neighborhoods by chapter or paragraph boundaries. Additionally, it may be interesting to use the absolute sentiment ratings instead of normalizing them to the average sentiment rating of the story. Seeing that some books are particularly sad or positive overall would be useful to the reader.

Furthermore, using a non-binary set of emotion dimensions might provide a more detailed view of the emotional content of the books. There are many dimensional emotion models that we could use for this purpose in addition to hypothesis-free topic modeling.

Finally, some aspects of the visualization could be improved from an aesthetic standpoint. For example, instead of using an instruction box, we could incorporate the legends into each aspect of visualization.

REFERENCES

- Bird, S., Loper, E., and Klein E. *Natural Language Processing with Python*. O'Reilly Media Inc. (2009).
- Bostock, M., Carter, M., Hinton, E., Horyn, C., and Wilson, E. "Front Row to Fashion Week." *nytimes.com*. New York Times, 12 Sept. 2013. Web. 13 Dec. 2013. <http://www.nytimes.com/newsgraphics/2013/09/13/fashion-week-editors-picks>

Cheng, N., Chandramouli, R., and Subbalakshmi, K.P. "Author gender identification from text." *Digital Investigation* 8(1) (2009): 78-88.

Cheng, N., Chen, X., Chandramouli, R., and Subbalakshmi, K.P. "Gender identification from e-mails." *Computational Intelligence and Data Mining (CIDM'09)*. Nashville, TN, USA. (2009).

Clark, J.. "Twitter StreamGraphs." *Twitter StreamGraphs*. 2009. Web. 13 Dec. 2013. <http://www.neoformix.com/Projects/TwitterStreamGraphs/view.php>

Eisenstein, J. and Barzilay, R. "Bayesian Unsupervised Topic Segmentation." In *Empirical Methods in Natural Language Processing (EMNLP)* (2008) 334-343.

Hearst, M. "Chapter 11. Information Visualization for Text Analysis", *Search User Interfaces*, Cambridge University Press, 2009.

Hu M. and Liu B. "Mining and summarizing customer reviews". *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery Data Mining (KDD-2004, full paper)*, Seattle, WA, USA (2004).

Hurst, Matthew. "Visualizing Jane Austen Part 2." 'Data Mining: Text Mining, Visualization and Social Media. Matthew Hurst, 23 Aug. 2011. Web. 13 Dec. 2013. http://datamining.typepad.com/data_mining/2011/08/visualizing-jane-austen-part-2.html

Paley, B. "TextArc Reading Alice's Adventures in Wonderland." *TextArc Reading Alice's Adventures in Wonderland*. TextArc, 1991. Web. 13 Dec. 2013. <http://www.textarc.org/Alice2.html>.

Posavec, S. "About This Project : Stefanie Posavec." *About This Project : Stefanie Posavec*. Stefanie Posavec, 2010. Web. 13 Dec. 2013. <http://www.itsbeenreal.co.uk/index.php/?wwwwords/about-this-project>.

Socher R., Perelygin A., Wu J., Chuang J., Manning C., Ng A., and Potts C. "Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank". *Conference on Empirical Methods in Natural Language Processing (EMNLP 2013, Oral)*. Seattle, WA, USA (2013).

Sudhahar, S., Franzosi, R., and Cristianini, N. "Automating Quantitative Narrative Analysis of News Data". *The Journal of Machine Learning Research (JMLR): Workshop and Conference Proceedings* 17 (2011) 63-71.

"Applying Sentiment Analysis to the Bible." *OpenBibleinfo*. OpenBibleinfo, 10 Oct. 2011. Web. 13 Dec. 2013. <http://www.openbible.info/blog/2011/10/applying-sentiment-analysis-to-the-bible>

"BookStats 2012 HIGHLIGHTS Annual Survey Captures

Size and Scope of US Publishing Industry." *BookStats*. Association of American Publishers, n.d. Web. 13 Dec. 2013. <http://bookstats.org/bookstats-2012.php>

"Bowker - Bowker Data." *Bowker - Bowker Data*. Bowker Market Research, n.d. Web. 13 Dec. 2013. http://www.bowker.com/en-US/aboutus/about_bowkerdata.shtml

"MovieBarcode." *Movie Barcode*. MovieBarcode, 2011. Web. 13 Dec. 2013. <http://moviebarcode.tumblr.com>

"Wordle - Obama's Nomination Acceptance Speech." *Wordle - Obama's Nomination Acceptance Speech*. wordle.net, n.d. Web. 13 Dec. 2013. http://www.wordle.net/show/wrdl/148852/Obama's_Nomination_Acceptance_Speech

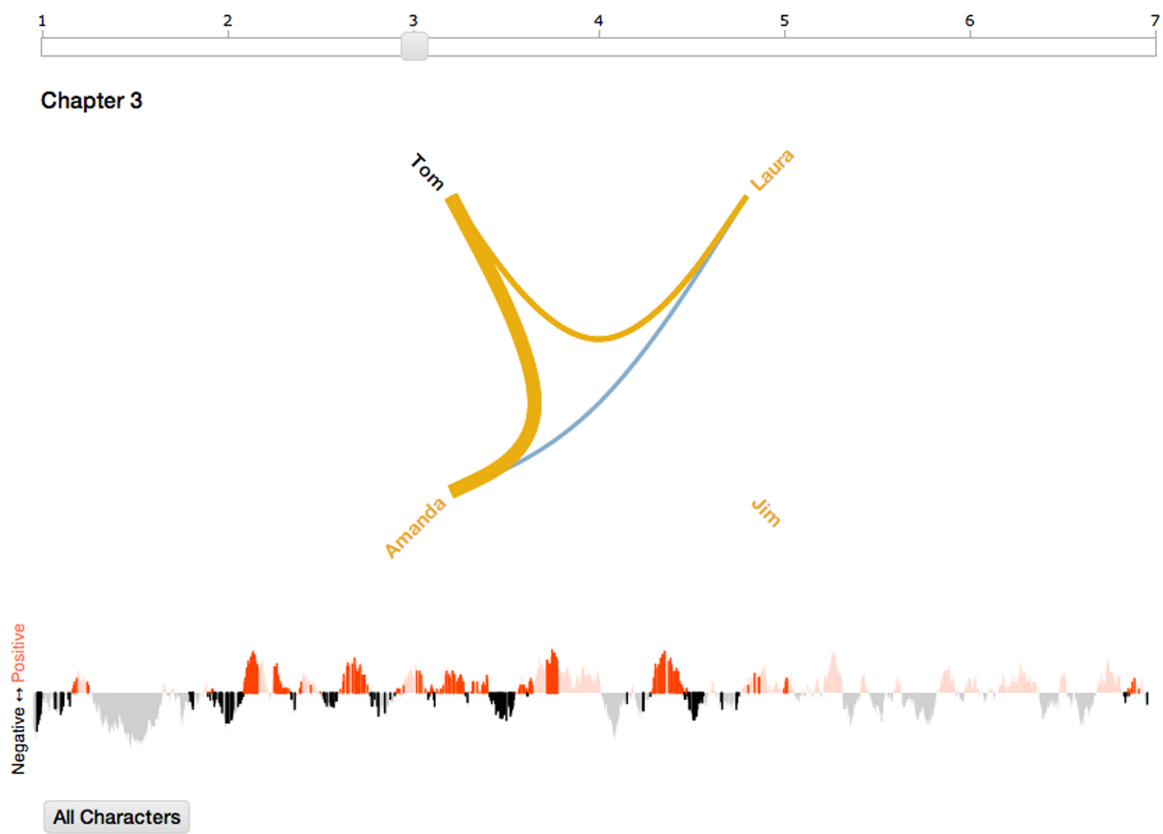


Fig. 10: Visualization of *The Glass Menagerie* by Tennessee Williams

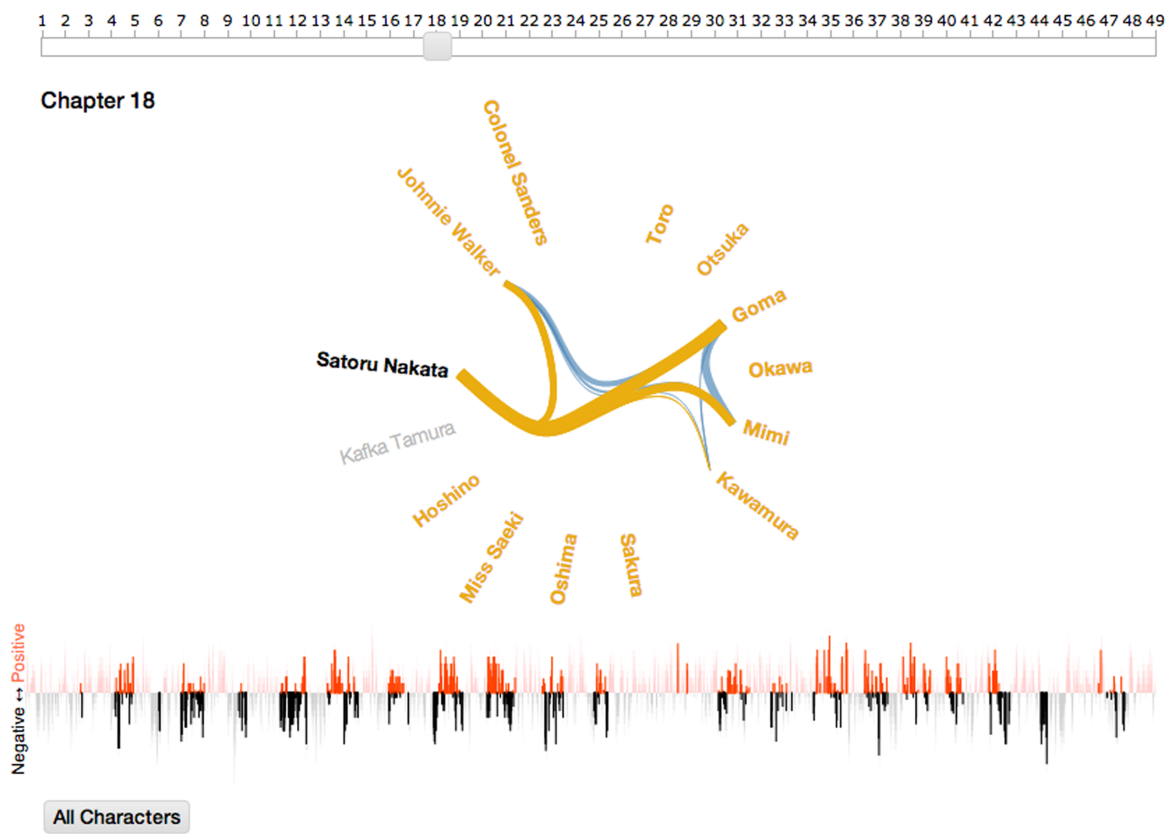


Fig. 11: Visualization of *Kafka on the Shore* by Haruki Murakami

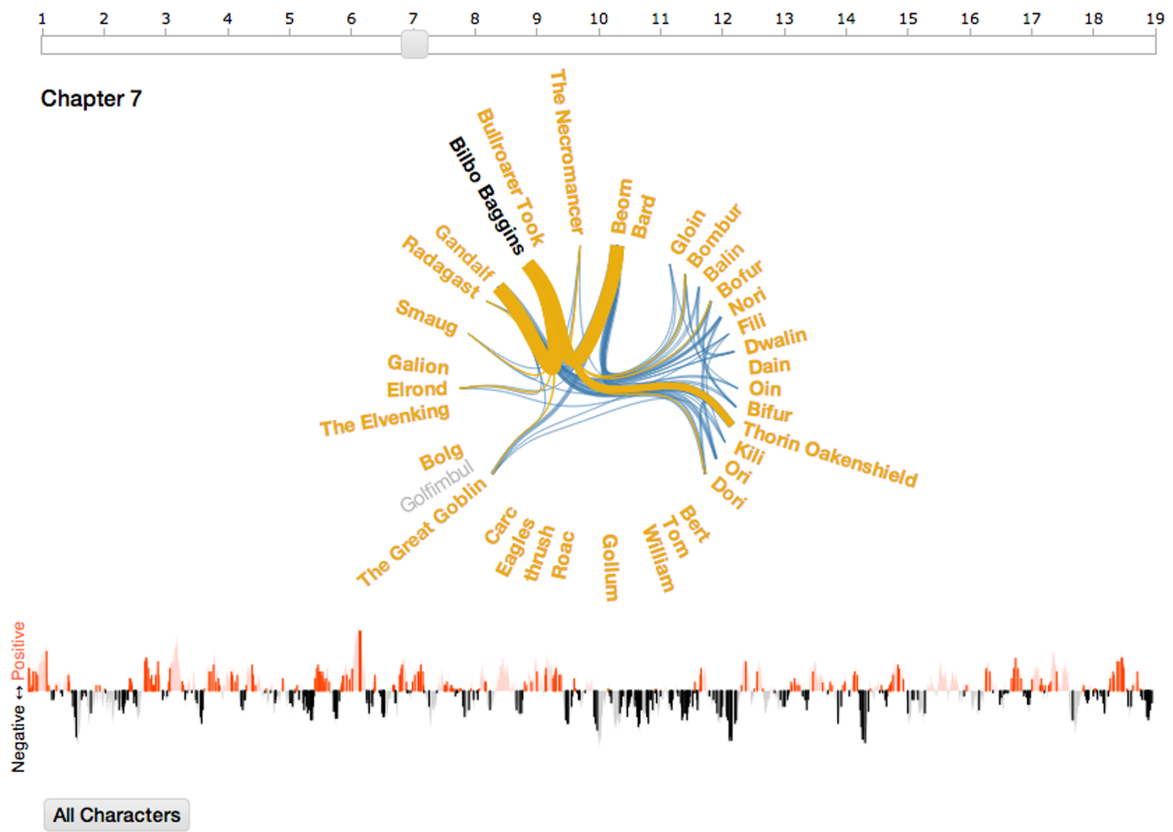


Fig. 12: Visualization of *The Hobbit* by J.R.R. Tolkien