Visual Analysis of Character and Plot Information Extracted from Narrative Text

Markus John[1(✉)], Steffen Lohmann[2], Steffen Koch[1], Michael Wörner[1], and Thomas Ertl[1]

[1] Institute for Visualization and Interactive Systems (VIS), University of Stuttgart, Universitätsstraße 38, 70569 Stuttgart, Germany
markus.john@vis.uni-stuttgart.de

[2] Fraunhofer Institute for Intelligent Analysis and Information Systems (IAIS), Schloss Birlinghoven, 53757 Sankt Augustin, Germany

Abstract. The study of novels and the analysis of their plot, characters and other information entities are complex and time-consuming tasks in literary science. The digitization of literature and the proliferation of electronic books provide new opportunities to support these tasks with visual abstractions. Methods from the field of computational linguistics can be used to automatically extract entities and their relations from digitized novels. However, these methods have known limitations, especially when applied to narrative text that does often not follow a common schema but can have various forms. Visualizations can address the limitations by providing visual clues to show the uncertainty of the extracted information, so that literary scholars get a better idea of the accuracy of the methods. In addition, interaction can be used to let users control and adapt the extraction and visualization methods according to their needs. This paper presents ViTA, a web-based approach that combines automatic analysis methods with effective visualization techniques. Different views on the extracted entities are provided and relations between them across the plot are indicated. Two usage scenarios show successful applications of the approach and demonstrate its benefits and limitations. Furthermore, the paper discusses how uncertainty might be represented in the different views and how users can be enabled to adapt the automatic methods.

Keywords: Text visualization · Visual text analytics · Digital humanities · Distant reading · Narrative text · Uncertainty visualization · NLP

1 Introduction

Common tasks in literary science are studying novels and analyzing their plot, characters and other entities. Literary scholars are interested in getting an overview of the plot and its characters, the relationships between them and their evolution during the plot [37]. The digitization of literature and the proliferation of electronic books (*ebooks*) provide new means to support these tasks with visual abstractions that are automatically generated from ebooks.

J. Braz et al. (Eds.): VISIGRAPP 2016, CCIS 693, pp. 220–241, 2017.
DOI: 10.1007/978-3-319-64870-5_11

Traditionally, literary scholars read and analyze novels in a sequential way by using so-called *close reading*. In contrast to this, Moretti introduced the idea of *distant reading* [32]. Instead of carefully reading and analyzing a literary work, distant reading abstracts the text by providing visualizations, such as graphs that depict the genre change of historical novels, maps to represent geographical aspects of the plot, or trees to classify various types of detective stories [21]. These visual abstractions can convey useful information and assist in exploring and understanding complex relationships, verifying hypotheses as well as forming new research ideas.

In order to provide visual abstractions for literary works, a combination of automatic methods and interactive visualization techniques is required. When dealing with ebooks, natural language processing (NLP) methods are the first choice for automatic analysis. Using these methods, entities such as characters and places can be extracted from ebooks. This enables the development of visual abstractions that allow to explore these entities and their relationships in more detail.

However, the combination of automatic NLP methods and visual representations often involve a certain degree of uncertainty. The NLP methods are normally trained on specific text corpora, such as newspaper or journal article texts, and do not provide entirely correct results for texts that differ from the training corpus thematically, chronologically or stylistically. Since the visual abstractions are derived through NLP methods, the uncertainty, or rather the errors, are reflected in them.

However, inaccurate methods can facilitate, for example, the work of a human annotator by highlighting and providing interactive adaptation possibilities of the extracted text characteristics. Visual clues could indicate uncertain annotations, which are good candidates for manual annotation, and users can correct or confirm these classifications and trigger a retraining to improve the NLP method, thereby implementing a feedback loop as discussed by Sacha et al. [41]. Furthermore, visual clues can assist literary scholars in getting a better idea about the accuracy of the NLP methods and ease text exploration and decision making.

This paper is an extended version of a work presented at IVAPP 2016 [24]. It presents ViTA, a web-based approach that aims to provide literary scholars with visual abstractions to facilitate character analysis in novels. It utilizes automatic named entity extraction and visualizes relationships between characters and places based on their co-occurrence, also considering temporal aspects. The basic idea of ViTA is to highlight patterns, such as specific characters and places or groups of characters that interact with each other at certain places over time. By offering several views, such patterns are made easily recognizable with the approach and provide the starting point for a deeper analysis. This can result in a better understanding of the plot, in particular related to the characters and their relationships.

The main contributions of this work are: (1) A web-based approach that offers a wide range of interactive features to facilitate character analysis in narrative texts. (2) Several visual abstractions that provide aggregated and interrelated information and allow for an interactive navigation to the corresponding passages in the text. (3) A discussion about the importance of uncertainty visualizations to support the analysis and decision-making, and how users can be enabled to adapt automatic methods.

The rest of the paper is structured as follows: Sect. 2 summarizes related work, before the ViTA approach is detailed in Sect. 3. This is followed by two use cases demonstrating the applicability and usefulness of ViTA in Sect. 4. Section 5 provides a discussion of the approach and Sect. 6 concludes the paper with a summary and outlook on future work.

2 Related Work

Since our approach is concerned with the visual abstraction of text, we first summarize existing work in this area in Sect. 2.1. Next, we report on visual analytics attempts in the field of literary science and review systems that are most closely related to our approach in Sect. 2.2. Last, we discuss in Sect. 2.3 how uncertainty can be represented in visualizations and how users can steer and adapt automatic methods.

2.1 Visual Text Abstraction

Several techniques for visually summarizing and abstracting text documents have been developed over the last couple of years.

One compact visualization method related to our approach is *literature fingerprinting* [25], which uses a pixel-based technique that represents each text unit as a single pixel and visually groups them into higher level units. A similar technique is used in Seesoft [11], which has been designed as a visual fingerprint summarization of source code to graphically represent software statistics. The intention of Tilebars [17] is the visual representation of search results comparable to the fingerprint idea, while FeatureLens [9] also uses a pixel-based attempt to explore interesting text patterns and to find co-occurrences in texts.

Another popular technique to visually summarize text documents are *word clouds* [13,29]. They usually depict the most frequently used words of a text with the font size scaled according to the word frequencies. Word clouds enable literature scholars to get a first impression of the main terms and topics of a text and can provide a useful starting point for deeper analyses [19,47].

To visualize relational information of a text document, visualization techniques such as PhraseNets [46] and WordTrees [48] were suggested. They depict syntactic, lexical or hierarchical relationships that exist between the words of a text as node-link diagrams. In contrast, Oelke et al. [35] use an adjacency matrix to encode the development of relations between entities across a text document.

Inspired by Munroe's hand-drawn illustration "Movie Narrative Charts" [33], a new visual technique has emerged as so-called *storyline* or *plot view* visualization. It aims to portray the dynamic relationship between entities in a story over time. Tanahashi and Ma [45] as well as Liu et al. [28] propose design considerations and optimization approaches for generating aesthetically appealing storyline visualizations.

2.2 Visual Text Analytics

In recent years, several approaches for visual text analytics have been introduced in different domains. Examples can be found in social media [10], patent analysis [26] or opinion mining [34], among others.

There are also quite a number of visual analytics approaches in the field of literary science. Jänicke et al. [22] propose several techniques for the visualization and comparison of text that is reused in different documents in order to support literary scholars in discovering and exploring intertextual similarities. Abdul-Rahman et al. [1] present a rule-based solution for poetry visualization, allowing for high-level interactions with the end users in a closed loop. They use glyphs to encode phonetic units and visual links to show phonetic and semantic relationships. The VarifocalReader [27] supports literary scholars by combining distant and close reading and by enabling intra-document explorations through advanced navigation concepts. It integrates machine learning techniques, search mechanisms and several visual abstractions.

Oelke et al. [35] discuss the analysis of prose literature by using the aforementioned literature fingerprinting technique. Their approach visually abstracts implicit relationships between characters and encodes their development within the analyzed novel. However, it does not allow to directly work with the text resource. Vuillemot et al. [47] present the system POSvis, which extracts named entities from literary text and focuses on the exploration of networks of characters. POSvis offers multiple coordinated views, including word clouds and self-organizing graphs, equipped with filter methods to review the vocabulary of novels. While this is closely related to our work, we do not pursue the goal to review the vocabulary in the context of one or more entities filtered by part of speech. Instead, we aim to support the analysis of characters and their relationships in the storyline of a novel, based on named entity extraction and co-occurrence analyses.

Another system closely related to our work is Jigsaw [43], which has been designed to support analysts during foraging and sense-making activities in collections of textual reports and other sets of documents. It provides multiple coordinated views including lists, scatter plots, word clouds and graph visualizations that allow tracking entities and exploring their relationships across the document collections. Jigsaw has been primarily designed for *inter*-document analysis, whereas we are interested in *intra*-document analysis, i.e., we support the analysis of a single text document at a time. Apart from that, Jigsaw follows a rather generic approach that does not focus on fictional literature and the analysis of characters but provides general-purpose visualizations for different kinds of entities extracted from the documents.

2.3 Visualization of Uncertainty

The inclusion of information about the quality of data, over which people are reasoning, has been gaining attention across different domains, such as education [8], geographic and geologic mapping [14], or statistics [42]. For analysts, it is important to have access to information about the accuracy and reliability of data in order to correctly interpret the data and make informed decisions.

In particular in the domain of *geographic information systems*, a number of methods for uncertainty visualization have been proposed [30,36]. Griethe and Schumann [16] summarize different techniques to indicate uncertainty in data, including:

– *Visual variables*, such as color, size, position, angle, texture or transparency.
– *Extra objects*, such as labels, images or glyphs.
– *Animation*, such as speed, duration, motion blur, range or extent of motion.
– *Other human senses*, such as incorporation of acoustics, changes in pitch, volume, rhythm, vibration, or flashing textual messages.

As an example, Collins et al. [7] present a visualization intended to reveal the uncertainty and variability inherent in statistically-derived lattice structures. Their approach uses different visual variables, including transparency, color, and size, to expose the inherent uncertainty in statistical processing and thus helps analysts in making more informed decisions. Such visual clues can also help literary scholars in getting a better feeling of the accuracy of automatically extracted information, especially in distant reading when scholars are not deeply familiar with the texts. In addition, visual clues can assist scholars in finding and recognizing errors in visual abstractions faster.

There are a couple of attempts to combine automatic NLP and visualization methods to enhance text analysis tasks. However, only few works implement an interactive and visual approach for the user-steered adaptation of the NLP methods. Brown et al. [3] present a technique to develop similarity functions for machine learning models in a visual and interactive way. The technique does not require profound knowledge on the complex parameters of those models, but can also be controlled by users with little experience in machine learning. Similarly, Hu et al. [20] suggest a semantic interaction approach that allows users to adapt parameters by user interactions on spatializations of the underlying models. Endert et al. [12] also use such a semantic interaction approach to support sensemaking in document collections. It combines analytic interactions in a spatialization (e.g., document repositioning, text highlighting, search, annotations) with updates to the model responsible for generating the spatial layout. A related approach is presented by Heimerl et al. [18] that lets users interactively train a support vector machine for the classification of text documents.

3 Visual Text Analysis with ViTA

A large interest has grown in web-based systems for literature analysis that are easy to use and do not require any skills in computational linguistics [38]. Against this background, the ViTA approach has been implemented as a web application that is easily available to literary scholars and other user groups and does not require any installation on the user's side. The implementation is based on standard web technologies and can be run with a modern web browser supporting HTML5, SVG, CSS, and JavaScript.[1] It provides different visual abstractions representing specific characteristics of the analyzed narrative text and highlighting search results that illustrate the development of characters in a storyline.

The web application offers automatic methods for importing ebooks, extracting entity information and visualizing this information. The developed visualizations include word clouds, fingerprints of characters and places, a graph representation indicating connections between characters and a plot view that illustrates the relationships between characters and places in the story of a novel over time.

3.1 Text Processing

There are a variety of formats for the digital representation of novels. One widely used format is EPUB, which is a free and open standard that encodes structure and layout information besides the actual text of the novel. Many digital libraries, such as Project Gutenberg[2], offer ebooks in EPUB format or alternatively as plain text.

EPUB and plain text are also the two formats supported by our approach. It does not require the plain text to be structured in a specific way. However, if the structure of chapters or other metadata should be considered in the analysis, this information must be given in the text file. We therefore utilize some simple markup to structure ebooks provided in plain text, such as those by Project Gutenberg. The markup can be used to add chapter headings, line breaks and comments as well as other metadata (e.g., the title, author(s), publication date, publisher, edition or genre of a book) manually or automatically.

Once the ebook is loaded into the system, it is processed in a linguistic analysis pipeline, consisting of tokenization, sentence splitting and named-entity recognition. The ViTA implementation offers three different analysis tools that users can choose from: Stanford CoreNLP[3], OpenNLP[4] and ANNIE[5]. All three tools perform state-of-the-art NLP but use different techniques that each have their advantages and limitations. Depending on the use case and type of novel, users can select the NLP tool that is most suitable for the analysis. As this is often not clear from the start, they can also run the linguistic analysis several times with all three tools, compare the different outputs and choose the one that produces the best results. The current implementation only supports the processing of English texts; however, it can be extended to other languages if required.

[1] A video and demo of the web implementation are available at http://textvis.visualdataweb.org.

[2] http://www.gutenberg.org.

[3] http://nlp.stanford.edu/software/corenlp.shtml.

[4] http://opennlp.apache.org.

[5] https://gate.ac.uk/ie/annie.html.

The users can set several other parameters to configure the analysis and visualizations. Most importantly, they can control whether stop words are removed and whether person and place names starting with a lowercase letter should be considered. These configuration parameters are only shown on demand and are intended for experienced users. By default, ANNIE is used for NLP processing and unlikely character and place names are removed, as these settings produced the best results for most of the novels we tested.

3.2 Overview Page

After an ebook has been linguistically analyzed, an overview page is shown, listing metadata about the book and providing links to the visual abstractions. As an example, the overview page of the classic adventure novel "Around the World in 80 Days" by Jules Verne is shown in Fig. 1.

The metadata is listed in the middle of the screen (Fig. 1ⓑ). Some of it (e.g., the title, author, and release date) is directly taken from the text source (if provided), while the computation of other metadata (e.g., the number of words and chapters) requires some basic text analysis. Yet other metadata, such as the main characters listed on the overview page, can only be determined by using advanced text analysis, in this case named-entity recognition. This advanced text analysis is computationally expensive and can take some time depending on the size of the novel. For instance, the advanced analysis of the novel "Around the World in 80 Days" with the NLP tool ANNIE requires around 50 s on the current server that hosts the demo application (Intel Core i7-4930K with 3.4 GHz and 4 GB RAM), whereas the basic analysis is completed in less than 5 s.

As a general strategy, we therefore decided to show the results of each analysis step whenever they are available; for example, the web application already shows the results of the basic analysis even though the advanced analysis is still running. This strategy also applies to the visualizations provided by our

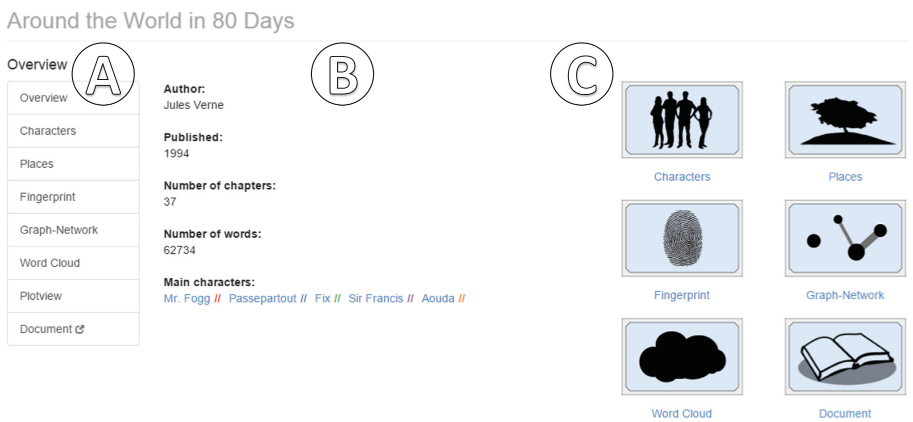

Fig. 1. Overview page listing metadata about the novel "Around the World in 80 Days".

approach. For instance, while the basic word cloud visualization can quickly be generated, the advanced word cloud as well as most other visualizations require more sophisticated NLP processing and can therefore not be shown before this processing is completed.

The interactive visualizations are available via the menu on the left (Fig. 1ⓐ), or via the icons on the right which are shown on the overview page (Fig. 1ⓒ). The fingerprint visualizations of the main characters can be directly opened from the overview page by clicking on the character links. The main characters are determined by counting the occurrences of all characters in the novel, with those that appear most often are assumed to be the protagonists. This simple measure worked surprisingly well for the novels we tested, in particular, since we also consider variations of the character names as detailed in Sect. 3.4.

The application assigns a unique color to each of the main characters. This color is shown on the overview page and constantly used for that character on all pages and in all visualizations. We created two color schemes, one for users with color vision deficiencies, consisting of four distinct colors determined by using the ColorBrewer 2.0[6], and the other for users with normal vision consisting of seven distinct colors using the categorical color scheme of D3[7]. If there are more than seven characters in a novel, the rest of them are shown in a gray color.

3.3 Characters and Places

To get an overview of the extracted characters and places, users can open either the characters or places view. Initially, the most frequently occurring entity is preselected in both views, complemented by a list of all extracted entities (Fig. 2ⓒ), where it is possible to search for and switch between entities.

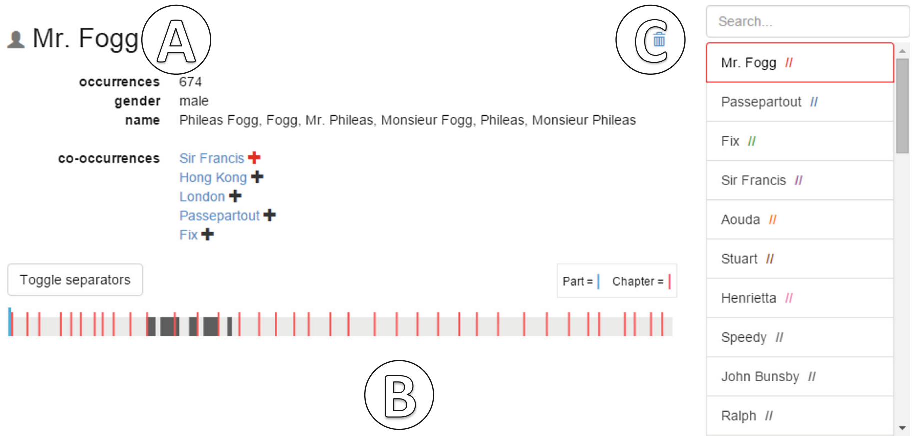

Fig. 2. Character view listing information about the character *Mr. Fogg* extracted from the novel.

[6] http://colorbrewer2.org.
[7] https://github.com/mbostock/d3/wiki/Ordinal-Scales#categorical-colors.

Fig. 3. Combined fingerprint visualization of two selected characters: *Passepartout* and *Fix*.

For each selected entity, a profile is provided, containing information about its occurrences, the detected gender and alternative names (for characters), and listing other entities that co-occur most frequently with that entity (Fig. 2ⓐ). In addition, the fingerprint visualization of the selected entity is shown (Fig. 2ⓑ).

3.4 Fingerprint Visualization

After the users obtained a first overview, they can further analyze the characters and places with the fingerprint visualization. It shows the temporal distribution of the entity occurrences in the novel. Blue and red bars represent parts and chapters of a book. Black blocks depict text segments where the entities occur. By clicking on the *toggle separators* button, the part and chapter bars can be hidden to get the 'plain' fingerprint.

Longer blocks indicate that entities occur often in that segment of the book, while shorter blocks appear when a character or place is only briefly mentioned. However, a longer block does not necessarily mean that entities are mentioned in every consecutive sentence, but it is sufficient if they are mentioned every few sentences.

Users can highlight a block segment by hovering over it to determine the respective chapter, shown in a tooltip. By clicking on it, they can jump to the corresponding text passage in the novel, which is opened in the text view. In that view, all occurrences are highlighted with the assigned specific color of the entity. This supports users in finding and analyzing text passages faster.

In addition, they can select multiple characters or places to get a combined fingerprint, as depicted in Fig. 3. The example shows the conjunction of two selected characters (*Passepartout* and *Fix*). This way, users can easily determine text passages where selected characters co-occur.

By default, we define that entities co-occur if they appear together in at least one sentence, which is a common co-occurrence metric. However, other co-occurrence metrics (range of words, etc.) are also possible and can be set as internal system parameters if required.

3.5 Character Network

To further investigate the character relations, users can switch to the character network view. This view contains a force-directed graph visualization that

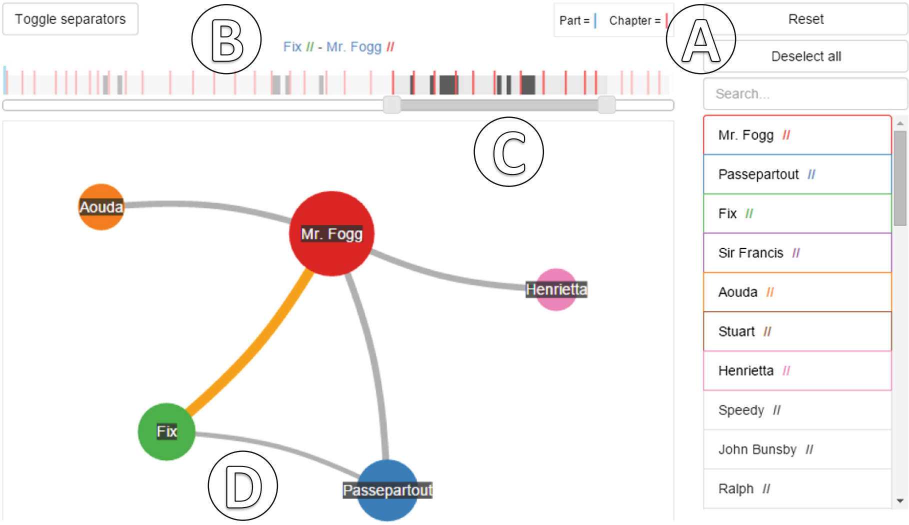

Fig. 4. Character network visualization representing the co-occurrence graph of a selected text segment. This example represents connections between the main characters of the novel "Around the World in 80 Days". The nodes of the graph represent the characters and the edges the number of sentences in which they co-occur.

represents connections between characters, as depicted in Fig. 4. The nodes of the graph represent the characters and the edges the number of sentences in which each pair of characters co-occurs. The node size and edge thickness are scaled proportionally to the characters' individual and co-occurrence frequency respectively. This helps to get a quick overview of the main characters and their connections.

The character network is once again complemented by a fingerprint visualization (Fig. 4ⓑ), indicating where the characters occur in the novel, and with a range slider (Fig. 4ⓒ) that lets users limit the view to a certain range of the novel (e.g., a single chapter). This has the advantage that users are not only enabled to analyze the overall structure of a novel but also the course of the relationships between characters, at least on the level of character co-occurrences. Users can select an edge in the graph (Fig. 4ⓓ) to display the co-occurrences of two related characters in the combined fingerprint visualization.

Initially, up to seven characters are preselected for the graph visualization, based on their occurrence frequency. The list of characters is displayed in that view again (Fig. 4ⓐ), so that users can search for characters and decide which ones are shown in the graph visualization by selecting or deselecting them in the list. That way, the graph visualization can be dynamically adapted according to the goals of the user. It can also be panned, zoomed, and rearranged to further support the analysis.

3.6 Word Cloud

If users are interested in terms that appear together with character mentions or in getting a first impression of the content of a book, they can switch to the word cloud view. Word clouds are commonly used by literary scholars, as they are considered easy to understand despite all their limitations [5,31]. The font size of the words is scaled proportionally to their occurrence frequency to visually indicate this value and put it in relation to other words. The exact word frequencies can be viewed in tooltips on demand.

The user can switch between a global word cloud representing the entire novel and local word clouds for the individual characters. The latter show the words that co-occur most often with the characters, as depicted in Fig. 5 for *Mr. Fogg*. This gives users some flexibility in their analysis, by providing a visually appealing overview of the novel or a novel character as well as supporting the discovery of new ideas and hypotheses.

3.7 Plot View

Finally, users can switch to the plot view to get a better idea of the dynamic relationships between characters. It reuses and extends an implementation of the University of Waterloo[8], which takes annotation data and automatically generates a storyline visualization in the spirit of the aforementioned "Movie Narrative Charts" [33].

Our approach adapts the visualization and displays the ten most frequently occurring characters as lines and every chapter as a node (cf. Fig. 6ⓐ). The horizontal axis represents the plot of the novel and the vertical grouping of lines indicates which characters co-occur in the chapters. If two or more lines share a node, this means that the corresponding characters co-occur frequently in that chapter. When hovering over a node, a tooltip lists the characters and places which co-occur in that chapter. Hovering over a line highlights the whole line as well as the name of the corresponding character.

The plot view supports users in getting a rough idea of the course of the storyline. It allows to quickly identify when and where characters come together or go separate ways and whether groups of characters exist. In the plot view of Fig. 6, one can see, for example, that *Mr. Fogg* and *Fix* interact for the first time in chapter seven (Fig. 6ⓑ), while *Aouda* and *John Bunsby* enter the plot not before the middle of the novel (Fig. 6ⓒ).

3.8 Text View

To support literary scholars in their common workflow and allow for close reading, we also provide a text view where they can directly work with the text. Recognized chapters are listed as hyperlinks on the left (Fig. 7ⓐ), while the text is presented on the right (Fig. 7ⓑ). The focused chapter is emphasized with bold

[8] http://csclub.uwaterloo.ca/~n2iskand/?page_id=13/.

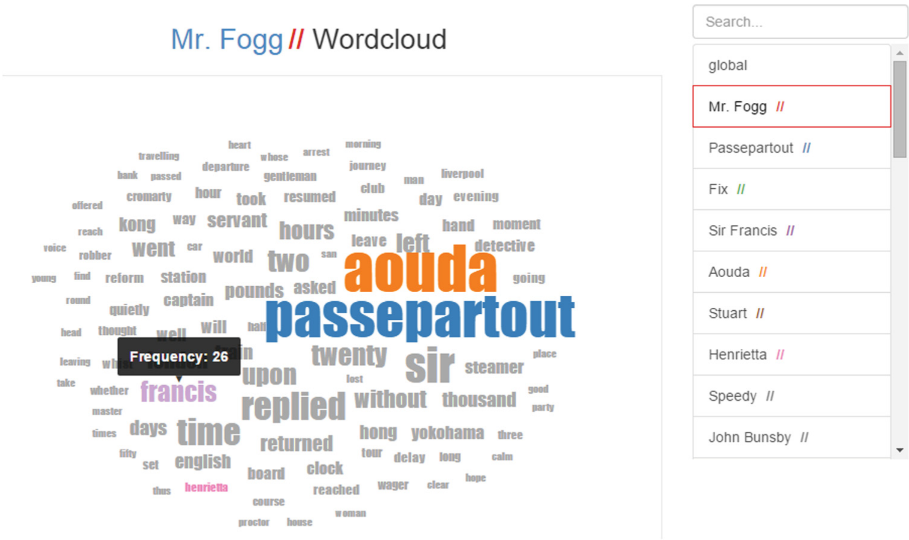

Fig. 5. Word cloud for the character *Mr. Fogg.*

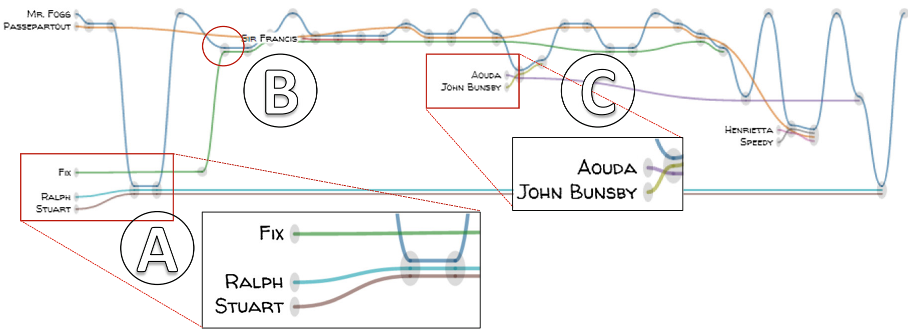

Fig. 6. Plot view of Jules Verne's novel "Around the World in 80 Days".

type. After clicking on a chapter, the text view jumps to the beginning of that chapter. If the user reached the text view from the fingerprint, graph network or word cloud view, the selected entity or entities are highlighted in the assigned color. Furthermore, there is a possibility for searching any other word or text passage as well as to reset the highlighting with a reset button.

The text view displays a vertical fingerprint next to its scrollbar. The idea is to provide both a visual representation of the distribution of entities and the possibility to inspect a text passage in detail, in order to support both distant and close reading. When hovering over the fingerprint blocks, the corresponding text passage is displayed in a tooltip (Fig. 7ⓑ), and after clicking on one, the text view jumps to the corresponding position. Additionally, the literary scholars

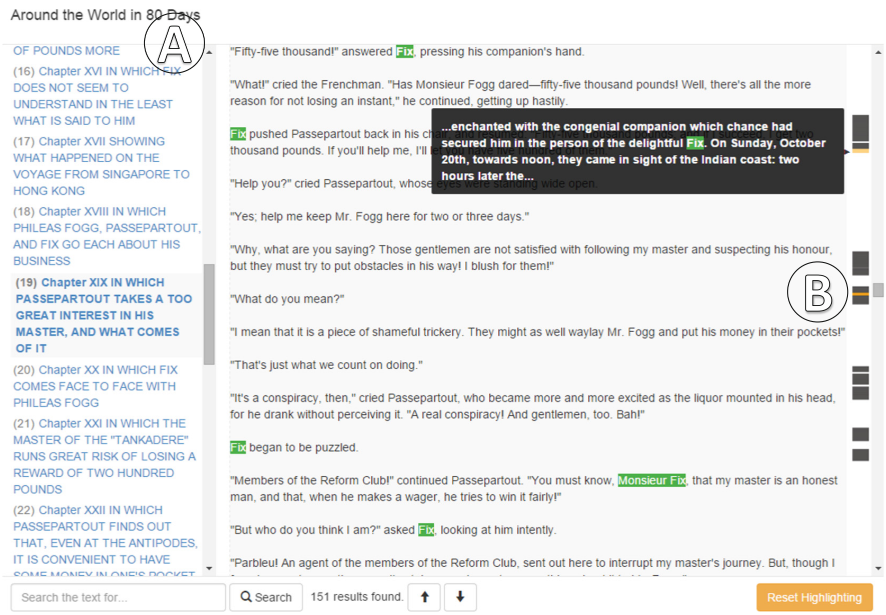

Fig. 7. Text view with selected character *Fix*.

can navigate to the next or previous occurrence of an entity by clicking on the up and down arrow buttons.

4 Usage Scenarios

In the following, we present two usage scenarios that demonstrate the applicability and usefulness of the ViTA approach. In the ePoetics project[9], we developed methods to support the analysis of German poetics – a form of early scholarly works on literature that formed the foundation of modern literature science. During our collaboration with colleagues from the literature department, we discussed the usefulness of direct access to text sources when using visual abstractions. This discussion inspired the development of an approach for analyzing narrative texts in a similar way, but taking into account the specific aspects of the plot and the social network of characters described in the text.

Although this was not the primary focus of the project, our colleagues emphasized its usefulness for their own work and for teaching. We therefore decided to set up a corresponding student project in oder to find out whether creating such visual abstractions from text mining results would be possible using off-the-shelf NLP techniques and tested the applicability on well-known novels.

For the usage scenarios, we selected a modern and an old English novel for analysis by a fictitious literary scholar. She has previous knowledge about the

[9] http://www.epoetics.de.

novels, since she read them some time ago, and is now trying to retrace the storyline and to detect the most important characters and events with our ViTA approach.

4.1 Analysis of the Novel "Harry Potter and the Half-Blood Prince"

In our first usage scenario, we present an analysis of the novel "Harry Potter and the Half-Blood Prince" by J. K. Rowling. It is the sixth and penultimate novel in the Harry Potter series and was published in 2009. The series chronicles the adventures of the young wizard Harry Potter and his quest to defeat the dark wizard Lord Voldemort, who strives to rid the wizarding world of Muggle (non-magical) heritage.

In a first step, the literary scholar explores and analyzes the character and network view. That way, she gets a quick overview of the main characters and their relationships. During the analysis, she encounters the name Slughorn and is surprised because she cannot remember him. To find out more about Slughorn, she selects the name in the character view (Fig. 8ⓐ) and uses the fingerprint visualization to navigate directly to its first occurrence, which is opened in the text view.

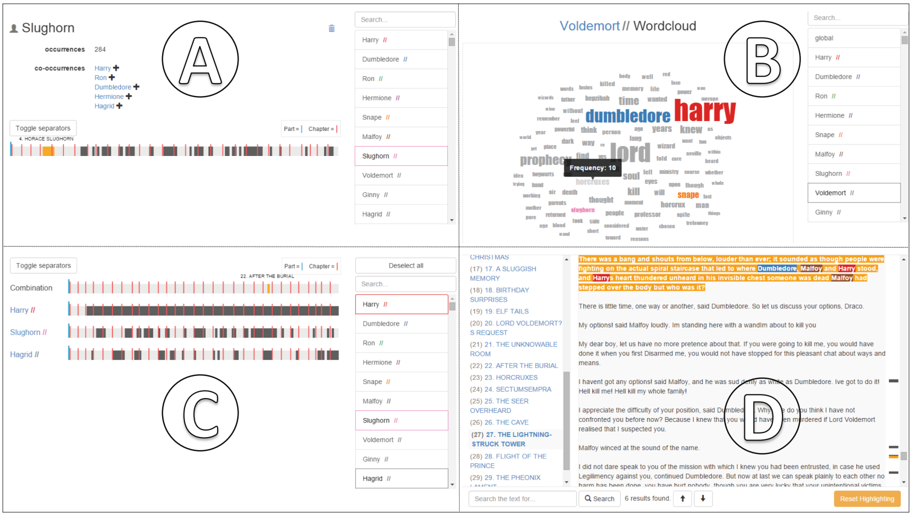

Fig. 8. Some of the visualizations used in the analysis of the novel "Harry Potter and the Half-Blood Prince": character view ⓐ, word cloud ⓑ, fingerprint visualization ⓒ and text view ⓓ.

While reading some paragraphs in the text view, she finds out that Dumbledore, the headmaster of the wizarding school Hogwarts, convinced Slughorn to return as potions teacher. Afterwards, she vaguely remembers that Dumbledore tasked Harry Potter with retrieving a memory from Slughorn that contains crucial information about Voldemort. To pursue this assumption, she switches to the word cloud view and explores the word clouds of Slughorn and Voldemort (Fig. 8ⓑ).

In the word cloud of Voldemort, she identifies the term Horcrux (an object in which a dark wizard has hidden a fragment of his soul for the purpose of attaining immortality) and remembers a conversation between Slughorn and Harry Potter at the home of Hagrid, the gamekeeper of Hogwarts.

In order to find the text passage, she uses the fingerprint view and selects the three characters Slughorn, Hagrid and Harry Potter. She determines that all three characters only co-occur at one text passage (Fig. 8ⓒ). Consequently, she jumps to that text passage and finds that Harry Potter succeeds in retrieving the memory, which shows Voldemort asking for information on creating Horcruxes.

In the following, she switches to the plot view since she is interested in examining the course of the storyline again. She still knows that Dumbledore dies at the end of the book and that Draco Malfoy, the son of one of Voldemort's followers, and Severus Snape, a professor at Hogwarts, are involved in his death. By analyzing the different chapters (nodes) and occurring characters (lines) in the plot view, she quickly recognizes the chapter of Dumbledore's death and navigates directly to the text of this chapter.

With the aid of the vertical fingerprints next to the text view's scroll bar, she can easily analyze the relevant text passages as depicted in Fig. 8ⓓ. She confirms her recollection that Draco Malfoy was chosen by Voldemort to kill Dumbledore. Furthermore, she finds out that Malfoy was unable to bring himself to do it and that Snape accomplished it with a deadly curse.

4.2 Analysis of the Novel "The Hobbit"

In the second usage scenario, our fictitious literary scholar analyzes the children's book "The Hobbit" by J. R. R. Tolkien. It was published in 1937 and is about Bilbo, a hobbit and the protagonist, and his adventures with dwarfs, elves, trolls and a dragon.

To reproduce the course of the novel, she starts her analysis on the plot view, as depicted in Fig. 9. By scanning the view, she gets a quick overview of the plot and remembers that Bilbo's adventure begins at his home with Gandalf, a wizard, and 13 dwarfs (Fig. 9ⓐ).

After jumping to the text and reading some passages in the text view, she remembers that they want to recover the treasure from Erebor (also known as the Lonely Mountain, former home to the greatest dwarf kingdom) and Bilbo is hired as their "burglar", since hobbits are small and unobtrusive.

Once she returns to the plot view, she notices the name of the creature Gollum (Fig. 9ⓑ) – originally a hobbit – and recalls that Bilbo wins a magical ring from him in a riddle war. However, she is unsure which further role Gollum plays in the plot and whether he co-occurs with other characters. To gain insights into this question, she switches to the graph view and immediately recognizes that

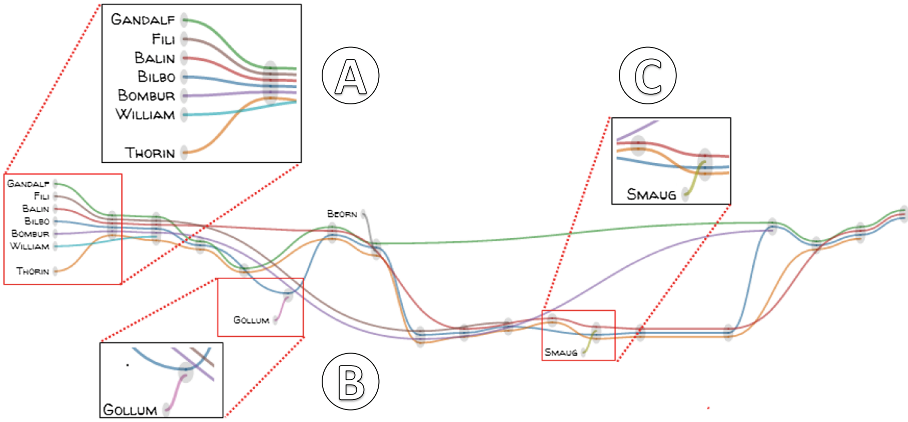

Fig. 9. Plot view showing some of the main characters of J. R. R. Tolkien's novel "The Hobbit".

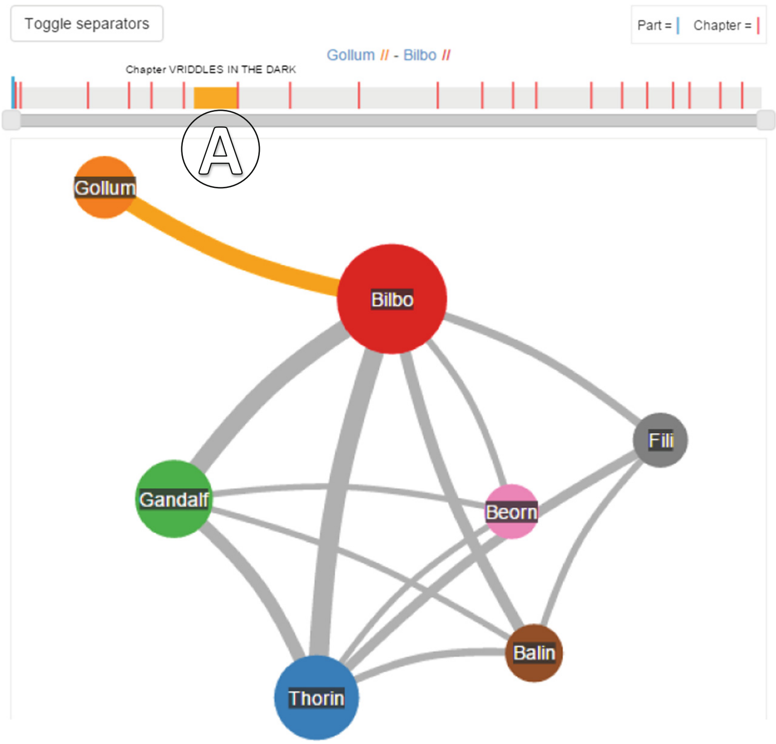

Fig. 10. Character network showing some of the main characters of "The Hobbit" and their co-occurrence frequencies. It can be observed that the character "Gollum" interacts only with "Bilbo" in the novel, while there seem to be strong ties between "Gandalf", "Thorin", "Bilbo" and the other main characters in the novel.

only a relationship between Gollum and Bilbo exist, as depicted in Fig. 10ⓐ. Subsequently, she activates the character view of Gollum to examine his fingerprint visualization, and thus confirms that he only appears at this point within the novel.

As the next step, she further analyzes the plot view and discovers that Smaug (Fig. 9ⓒ), a great fire dragon, enters the plot at the end of the novel. The literary scholar knows that Smaug invaded the dwarf kingdom of Erebor a long time ago and now guards the treasure. She vaguely remembers that Bilbo and the dwarfs are searching for a specific treasure object. To inspect this assumption, she switches to the text view and jumps to text passage where Smaug occurs for the first time in the plot. While reading, she finds out that the searched-for object is the Arkenstone, a great jewel.

In order to get more information on the Arkenstone, she uses the keyword search to highlight all occurrences as vertical fingerprints next to the text view's scroll bar. By analyzing the text passages, she determines, for example, that the Arkenstone is a heirloom of the dwarf kings and that, at the end of the novel, it is placed upon the chest of Thorin – the legitimate king of Erebor – within his tomb deep under the Lonely Mountain.

The usage scenarios show that our ViTA approach provides literary scholars with visual abstractions that facilitate character analysis in novels. Through the developed set of visual and automatic methods, we support them in generating and investigating hypotheses, confirming recollections and gaining insights.

5 Discussion and Future Work

It is important to note that we developed the ViTA approach to support *intra-document* analysis by visualizing the dynamic relationships between characters in a narrative text. ViTA is not intended to support a comparative *inter-document* analysis. However, we provide several visual abstractions that offer different views on the text to support literary scholars and other user groups in their analysis.

5.1 Incorporation of Uncertainty Information

The implemented visual abstractions are based on out-of-the-box NLP toolkits and methods, which have continuously improved and show strong robustness. Typically, such NLP methods are either rule-based or rely on statistical and machine learning approaches, which have been trained on available newspaper or journal article texts, since large training sets are most often available for these types of texts. As a consequence, NLP methods could be less effective when applied to very specific types of text, such as historical or fictional novels. This can results in uncertainties and errors in the visual abstractions. An uncritical interpretation of the provided visualizations can therefore cause confusion and misunderstandings.

Interactive visualizations can play an important role in addressing these problems. We intend to provide visual clues that communicate the quality of an automatic analysis. This could help literary scholars to interpret automatically generated results more accurately. Colors are particularly suited for the representation of such visual clues, for example, green could represent a very reliable result, whereas red could indicate high uncertainty. As mentioned before, labels are also suitable to communicate uncertainty information, for example, by adding a confidence value (e.g., given as percentage) for each extracted entity. However, labeling can be challenging as the uncertainty information must not obscure other labels or graphical elements in the visualization, and must be easy to spot by users at the same time. Therefore, labels would be particularly useful in the character and places views of ViTA, as there is sufficient space to represent uncertainty information next to the detected entities.

Another possibility to indicate the accuracy of an extraction method is to use glyph-based techniques like Chernoff faces [4]. However, these techniques have the drawback that they do not scale to a large number of entities and that users first need to learn how to read them. A more appropriate method to show the confidence value of automatic methods with the help of visual clues is color saturation. For example, light colors could indicate entities with a high uncertainty, which are good candidates for manual correction, while strong colors could indicate entities which are assumed to be accurate. Using saturation to communicate uncertainty is easy to understand and can also be easily integrated in the visualizations provided by ViTA. For example, color saturation can be added to nodes and lines to represent the accuracy of the automatically generated results in the graph network or plot view. Moreover, saturation might be used to display the accuracy of recognized terms in the word cloud view or of the extracted entities in the text view. The use of animations to represent uncertainty has to be handled with care, as animations are known to quickly confuse users with transient information and may result in perceptual overload.

5.2 Other Extensions and Improvements

We plan to provide better means to interactively control and adapt automatic NLP methods as they are used in ViTA. In this regard, visual clues can support users in finding uncertain extraction results, which are good candidates for manual correction. Users could then correct or confirm the extracted information one after the other or according to some predefined priority schemes. This could finally help to improve the automatic methods and analysis results, also when applied to similar texts.

The current implementation for the detection of co-occurrences between entities is based on named entity recognition. We determine that two entities co-occur if they both appear in the same sentence. This approach works well in many cases, however, it can be improved in various aspects. One option would be to let the literary scholars set the co-occurrences range variable, for example, increase the context to several sentences or a whole paragraph. Also, methods for coreference resolution might be integrated to detect alternative occurrences of the entities in the text and compute more precise frequency values [39,44]. However, existing coreference techniques are again mostly trained on non-fictional texts and have to be adapted (where appropriate) to the type of text and user needs. We also plan to take the semantics of the text into account to better investigate the relationships between entities. Once again existing approaches [15,40] could be integrated to provide the literary scholars with additional information for their analyses. In general, we aim at supporting a broader spectrum of natural language processing methods in the future.

Another topic for future work is the challenge of visual scalability when working with very long texts. Most of the visual abstractions used in ViTA scale quite well. However, the fingerprint visualization and the plot view become less useful with a growing text length. Since the available screen space remains constant, they are increasingly compressed until they finally overlap. To address this problem, some focus+context technique could be integrated, such as a fisheye distortion [2], or an overview+detail approach to present multiple views with different levels of abstractions [6].

Apparently, the presented ViTA approach can also be extended by additional visualizations. For example, views that visualize geographical information [23] might be integrated if the story of the text is based on real world places, such as the above analyzed novel "Around the World in 80 Days", or own cartographic material in the case of fictional places like in the novel "The Hobbit". Furthermore, the existing visualizations could be improved and extended. One such extension could be different graph layouts for the character network to get a better impression of the constellation of characters or to add more information, such as semantic relations between characters. Another possible extension could be to equip the plot view with automatically detected events (e.g., the "Battle of Helm's Deep" in case of the novel "The Lord of the Rings"), enabling users to better track and identify the course of the story.

6 Conclusion

In this work, we presented an easily accessible web-based approach for visualizing the relation of characters and places in a novel. The approach incorporates different natural language processing toolkits to extract named entities, and provides possibilities to set parameters for variable analyses. Furthermore, it provides several visual abstractions that support literary scholars in a flexible and comprehensive analysis of the novel characters. The approach facilitates distant reading, in particular, and provides a starting point for new ideas, hypotheses and further analyses. All views enable literary scholars to jump to the corresponding text passages and thus allow for working with the text directly. We presented an implementation of the approach and provided two usage scenarios that illustrate its applicability and usefulness.

Acknowledgements. We would like to thank our students Sanjeev Balakrishnan, Felix Do, Sebastian Frank, Paul Kuznecov, Vincent Link, Eduard Marbach, Jan Melcher, Christian Richter, Marc Weise, and Marvin Wyrich who implemented the approach in a student project. This work has partly been funded by the German Federal Ministry of Education and Research (BMBF) as part of the 'ePoetics' project and as part of the Center for Reflected Text Analysis 'CRETA' at University of Stuttgart.

References

1. Abdul-Rahman, A., Lein, J., Coles, K., Maguire, E., Meyer, M., Wynne, M., Johnson, C.R., Trefethen, A., Chen, M.: Rule-based visual mappings - with a case study on poetry visualization. Comput. Graph. Forum **32**(3pt4), 381–390 (2013)
2. Bederson, B.: Fisheye menus. In: Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology, UIST 2000, pp. 217–225. ACM (2000)
3. Brown, E., Liu, J., Brodley, C., Chang, R.: Dis-function: learning distance functions interactively. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, VAST 2012, pp. 83–92 (2012)
4. Chernoff, H.: The use of faces to represent points in k-dimensional space graphically. J. Am. Stat. Assoc. **68**(342), 361–368 (1973)
5. Clement, T., Plaisant, C., Vuillemot, R.: The story of one: humanity scholarship with visualization and text analysis. In: Proceedings of the Digital Humanities Conference, DH 2009. HCIL-2008-33 (2009)
6. Cockburn, A., Karlson, A., Bederson, B.: A review of overview+detail, zooming, and focus+context interfaces. ACM Comput. Surv. **41**(1), 1–31 (2009)
7. Collins, C., Carpendale, S., Penn, G.: Visualization of uncertainty in lattices to support decision-making. In: Proceedings of the 9th Joint Eurographics/IEEE VGTC Conference on Visualization, EuroVis 2007, pp. 51–58. Eurographics Association (2007)
8. Epp, C.D., Bull, S.: Uncertainty representation in visualizations of learning analytics for learners: current approaches and opportunities. IEEE Trans. Learn. Technol. **8**(3), 242–260 (2015)
9. Don, A., Zheleva, E., Gregory, M., Tarkan, S., Auvil, L., Clement, T., Shneiderman, B., Plaisant, C.: Discovering interesting usage patterns in text collections: integrating text mining with visualization. In: Proceedings of the 16th ACM Conference on Conference on Information and Knowledge Management, CIKM 2007, pp. 213–222. ACM (2007)
10. Dou, W., Wang, X., Skau, D., Ribarsky, W., Zhou, M.: Leadline: interactive visual analysis of text data through event identification and exploration. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, VAST 2012, pp. 93–102 (2012)
11. Eick, S., Steffen, J., Sumner, E.E.: Seesoft-a tool for visualizing line oriented software statistics. IEEE Trans. Softw. Eng. **18**(11), 957–968 (1992)
12. Endert, A., Fiaux, P., North, C.: Semantic interaction for sensemaking: inferring analytical reasoning for model steering. IEEE Trans. Vis. Comput. Graph. **18**(12), 2879–2888 (2012)
13. Feinberg, J.: Wordle. In: Beautiful Visualization, pp. 37–58. O'Reilly (2010)
14. Fisher, P.F.: Visualizing uncertainty in soil maps by animation. Cartogr.: Int. J. Geog. Inf. Geovisualization **30**(2–3), 20–27 (1993)
15. Gildea, D., Jurafsky, D.: Automatic labeling of semantic roles. Comput. Linguist. **28**(3), 245–288 (2002)

16. Griethe, H., Schumann, H.: The visualization of uncertain data: methods and problems. In: Proceedings of Simulation und Visualisierung 2006, SimVis 2006, pp. 143–156. SCS (2006)
17. Hearst, M.A.: Tilebars: visualization of term distribution information in full text information access. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 1995, pp. 59–66. ACM/Addison-Wesley (1995)
18. Heimerl, F., Koch, S., Bosch, H., Ertl, T.: Visual classifier training for text document retrieval. IEEE Trans. Vis. Comput. Graph. **18**(12), 2839–2848 (2012)
19. Heimerl, F., Lohmann, S., Lange, S., Ertl, T.: Word cloud explorer: text analytics based on word clouds. In: 47th Hawaii International Conference on System Sciences, HICSS 2014, pp. 1833–1842. IEEE (2014)
20. Hu, X., Bradel, L., Maiti, D., House, L., North, C.: Semantics of directly manipulating spatializations. IEEE Trans. Vis. Comput. Graph. **19**(12), 2052–2059 (2013)
21. Jänicke, S., Franzini, G., Cheema, M.F., Scheuermann, G.: On close and distant reading in digital humanities: a survey and future challenges. In: Eurographics Conference on Visualization - STARs, EuroVis 2015. The Eurographics Association (2015)
22. Jänicke, S., Geßner, A., Büchler, M., Scheuermann, G.: Visualizations for text re-use. In: Proceedings of the 5th International Conference on Information Visualization Theory and Applications, IVAPP 2014, pp. 59–70. Scitepress (2014)
23. Jänicke, S., Heine, C., Stockmann, R., Scheuermann, G.: Comparative visualization of geospatial-temporal data. In: Proceedings of the 3rd International Conference on Information Visualization Theory and Applications, IVAPP 2012, pp. 613–625. Scitepress (2012)
24. John, M., Lohmann, S., Koch, S., Wörner, M., Ertl, T.: Visual analytics for narrative text - visualizing characters and their relationships as extracted from novels. In: Proceedings of the 7th International Conference on Information Visualization Theory and Applications, IVAPP 2016, pp. 27–38. Scitepress (2016)
25. Keim, D., Oelke, D.: Literature fingerprinting: a new method for visual literary analysis. In: Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, VAST 2007, pp. 115–122 (2007)
26. Koch, S., Bosch, H., Giereth, M., Ertl, T.: Iterative integration of visual insights during scalable patent search and analysis. IEEE Trans. Vis. Comput. Graph. **17**(5), 557–569 (2011)
27. Koch, S., John, M., Wörner, M., Müller, A., Ertl, T.: VarifocalReader - in-depth visual analysis of large text documents. IEEE Trans. Vis. Comput. Graph. **20**(12), 1723–1732 (2014)
28. Liu, S., Wu, Y., Wei, E., Liu, M., Liu, Y.: Storyflow: tracking the evolution of stories. IEEE Trans. Vis. Comput. Graph. **19**(12), 2436–2445 (2013)
29. Lohmann, S., Ziegler, J., Tetzlaff, L.: Comparison of tag cloud layouts: task-related performance and visual exploration. In: Gross, T., Gulliksen, J., Kotzé, P., Oestreicher, L., Palanque, P., Prates, R.O., Winckler, M. (eds.) INTERACT 2009. LNCS, vol. 5726, pp. 392–404. Springer, Heidelberg (2009). doi:10.1007/978-3-642-03655-2_43
30. MacEachren, A.M., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M., Hetzler, E.: Visualizing geospatial information uncertainty: what we know and what we need to know. Cartogr. Geogr. Inf. Sci. **32**(3), 139–160 (2005)
31. McNaught, C., Lam, P.: Using wordle as a supplementary research tool. Qual. Rep. **15**(3), 630–643 (2010)
32. Moretti, F.: Graphs, Maps, Trees: Abstract Models for a Literary History. Verso (2005)

33. Munroe, R.: Movie Narrative Charts (2009). http://xkcd.com/657/
34. Oelke, D., Hao, M., Rohrdantz, C., Keim, D., Dayal, U., Haug, L., Janetzko, H.: Visual opinion analysis of customer feedback data. In: Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, VAST 2009, pp. 187–194 (2009)
35. Oelke, D., Kokkinakis, D., Keim, D.A.: Fingerprint matrices: uncovering the dynamics of social networks in prose literature. Comput. Graph. Forum **32**(3pt4), 371–380 (2013)
36. Pang, A.T., Wittenbrink, C.M., Lodha, S.K.: Approaches to uncertainty visualization. Vis. Comput. **13**(8), 370–390 (1997)
37. Phelan, J.: Reading People, Reading Plots: Character, Progression, and the Interpretation of Narrative. University of Chicago Press, Chicago (1989)
38. Plaisant, C., Rose, J., Yu, B., Auvil, L., Kirschenbaum, M.G., Smith, M.N., Clement, T., Lord, G.: Exploring erotics in Emily Dickinson's correspondence with text mining and visual interfaces. In: Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries, JCDL 2006, pp. 141–150. ACM (2006)
39. Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., Manning, C.: A multi-pass sieve for coreference resolution. In: Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP 2010, pp. 492–501. ACL (2010)
40. Ruiz-Casado, M., Alfonseca, E., Castells, P.: Automatising the learning of lexical patterns: an application to the enrichment of wordnet by extracting semantic relationships from Wikipedia. Data Knowl. Eng. **61**(3), 484–499 (2007)
41. Sacha, D., Stoffel, A., Stoffel, F., Kwon, B.C., Ellis, G., Keim, D.A.: Knowledge generation model for visual analytics. IEEE Trans. Vis. Comput. Graph. **20**(12), 1604–1613 (2014)
42. Skeels, M., Lee, B., Smith, G., Robertson, G.G.: Revealing uncertainty for information visualization. Inf. Vis. **9**(1), 70–81 (2010)
43. Stasko, J., Görg, C., Liu, Z.: Jigsaw: supporting investigative analysis through interactive visualization. Inf. Vis. **7**(2), 118–132 (2008)
44. Stoyanov, V., Cardie, C., Gilbert, N., Riloff, E., Buttler, D., Hysom, D.: Coreference resolution with reconcile. In: Proceedings of the ACL 2010 Conference Short Papers, pp. 156–161. ACL (2010)
45. Tanahashi, Y., Ma, K.-L.: Design considerations for optimizing storyline visualizations. IEEE Trans. Vis. Comput. Graph. **18**(12), 2679–2688 (2012)
46. Van Ham, F., Wattenberg, M., Viegas, F.: Mapping text with phrase nets. IEEE Trans. Vis. Comput. Graph. **15**(6), 1169–1176 (2009)
47. Vuillemot, R., Clement, T., Plaisant, C., Kumar, A.: What's being said near "Martha"? Exploring name entities in literary text collections. In: Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, VAST 2009, pp. 107–114 (2009)
48. Wattenberg, M., Viegas, F.: The word tree, an interactive visual concordance. IEEE Trans. Vis. Comput. Graph. **14**(6), 1221–1228 (2008)