

Visual Analytics for Narrative Text

Visualizing Characters and their Relationships as Extracted from Novels

Markus John, Steffen Lohmann, Steffen Koch, Michael Wörner and Thomas Ertl
Institute for Visualization and Interactive Systems, University of Stuttgart, Universitätsstraße 38, Stuttgart, Germany

Keywords: Text Visualization, Visual Text Analytics, Digital Humanities, Distant Reading.

Abstract: The study of novels and the analysis of their plot, characters and other entities are time-consuming and complex 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 fields of computational linguistics can be used to automatically extract entities and their relations from digitized novels, which can then be visualized to ease exploration and analysis tasks. This paper presents 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.

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 (Phelan, 1989). 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.

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* (Moretti, 2005). 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 (Jänicke et al., 2015). 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 methods are the first choice for automatic analy-

sis. Using these methods, entities such as characters and places can be extracted from the ebooks. This enables the development of visual abstractions that allow to explore these entities and their relationships in more detail. In particular, 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 (Plaisant et al., 2006).

This paper presents 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 co-occurrence analyses. The basic idea of the approach 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 our 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 novels. 2) Several interrelated visual abstractions that enable the interactive switching to the corresponding text sources in the document.

The rest of the paper is structured as follows: Section 2 summarizes related work before our approach

is detailed in Section 3. This is followed by two use cases demonstrating the applicability and usefulness of the approach in Section 4. Section 5 provides a discussion of the approach and Section 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. Next, we report on visual analytics attempts in the field of literary science and review the systems that are most closely related to our approach.

2.1 Visual Text Abstraction

Several techniques for visually abstracting and summarizing text documents have been developed over the last years. One compact visualization method related to our approach is *literature fingerprinting* (Keim and Oelke, 2007), which uses a pixel-based technique that represents each text unit as a single pixel and visually groups them into higher level units.

A related technique is used in Seesoft (Eick et al., 1992), which has been designed as a visual fingerprint summarization of source code to graphically represent software statistics. The intention of Tilebars (Hearst, 1995) is the visual representation of search results similar to the fingerprint idea, while Feature-Lens (Don et al., 2007) 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 are word clouds (Viégas and Wattenberg, 2008; Clement et al., 2009). 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 (Heimerl et al., 2014).

To visualize relational information of a text document, approaches such as WordTrees (Wattenberg and Viegas, 2008) and PhraseNets (Van Ham et al., 2009) were suggested. They depict either syntactic, lexical or hierarchical relationships that exist between the words of a text as node-link diagrams. In contrast, (Oelke et al., 2013) 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" (Munroe, 2009), a new visual technique has emerged as so-called storyline or plot view visualization (Liu et al., 2013). It aims to

portray the dynamic relationship between entities in a story over time. (Tanahashi and Ma, 2012) as well as (Liu et al., 2013) propose design considerations and an efficient optimization approach 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 (Dou et al., 2012), opinion mining (Oelke et al., 2009) or patent analysis (Koch et al., 2011), among others.

There are also quite a number of visual analytics approaches in the field of literary science. (Jänicke et al., 2014) 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., 2013) 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 (Koch et al., 2014) 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., 2013) discuss the analysis of prose literature by using the aforementioned literature fingerprinting technique. Their approach visually abstracts implicit relationships between characters and, at the same time, encodes their development within the analyzed novel. However, it does not allow to directly work with the text resource.

(Vuillemot et al., 2009) 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 (Stasko et al., 2008), 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 designed primarily 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.

3 VISUAL ANALYSIS

Our approach has been implemented as a web application that is easily available to literary scholars and other user groups, as it 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 novel and highlighting search results that illustrate the development of characters in a storyline.

The web application offers automatic methods for importing novels, extracting characters and places and visualizing them. 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 a story 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 Gutenberg², 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 implementation offers three different analysis tools that users can choose from: Stanford CoreNLP³, OpenNLP⁴ and ANNIE⁵. All three tools perform state-of-the-art natural language processing (NLP) but use different techniques that each have their benefits 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.

The users can set several other parameters for the analysis and visualization. Most importantly, they can control whether stop words are removed and whether person and place names starting with a lowercase letter should be considered. However, these parameters are only shown on demand and are intended for the experienced user. By default, ANNIE is used for NLP processing and unlikely character and place names are removed, as these settings, in our experience, 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 Figure 16.

Some of the metadata (e.g., the title, author, and release date) is directly taken from the text source (if provided), while the determination 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,

¹A public demo of the web application is 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>

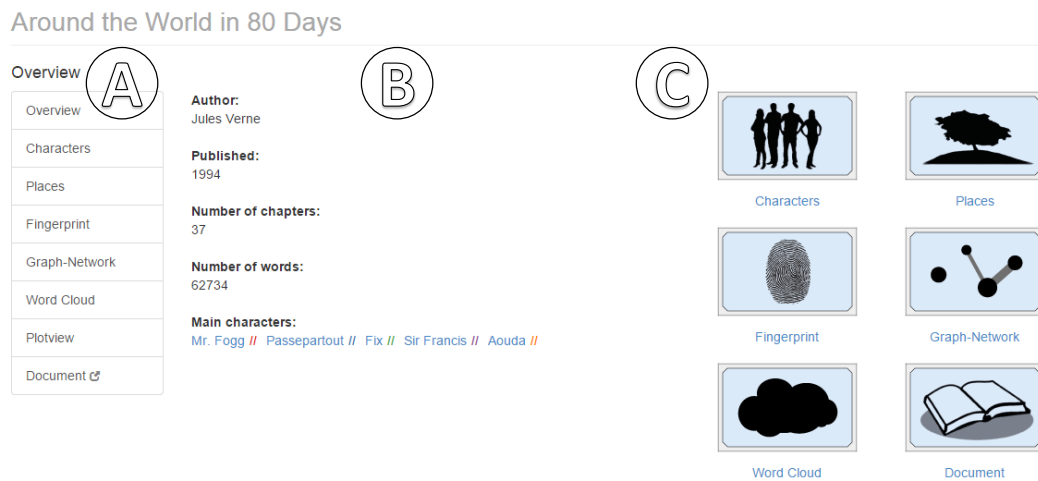


Figure 1: Overview page showing metadata about the novel “Around the World in 80 Days” by Jules Verne.

in this case named-entity recognition. This advanced text analysis is computationally complex 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 ANNIE requires around 50 seconds on the current server that hosts the demo application (CPU: Intel Core i7-4930K with 3.4 GHz, 4 GB RAM), whereas the basic analysis is completed in less than 5 seconds. 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 although the advanced analysis is still running.

This strategy also applies to the different visualizations provided by our approach. For instance, while a basic word cloud visualization can quickly be generated, the more 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 (Figure 1Ⓐ), or via the icons on the right on the overview page (Figure 1Ⓒ). The fingerprint visualizations of the main characters can directly be opened from the overview page by clicking on the corresponding links. The main characters are determined by counting their occurrences in the novel, with those that appear most often in the novel 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 Section 3.4.

The application assigns a unique color to each of the main characters. This color is shown in the overview page and consistently 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 2.0⁶, and the other for users with normal vision consisting of seven distinct colors using the categorical color scheme of D3⁷. 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 (Figure 2Ⓒ), where it is possible to search for and switch between entities.

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 (Figure 2Ⓐ). In addition, the fingerprint visualization of the selected entity is shown (Figure 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.

⁶<http://colorbrewer2.org>

⁷<https://github.com/mbostock/d3/wiki/Ordinal-Scales#categorical-colors>

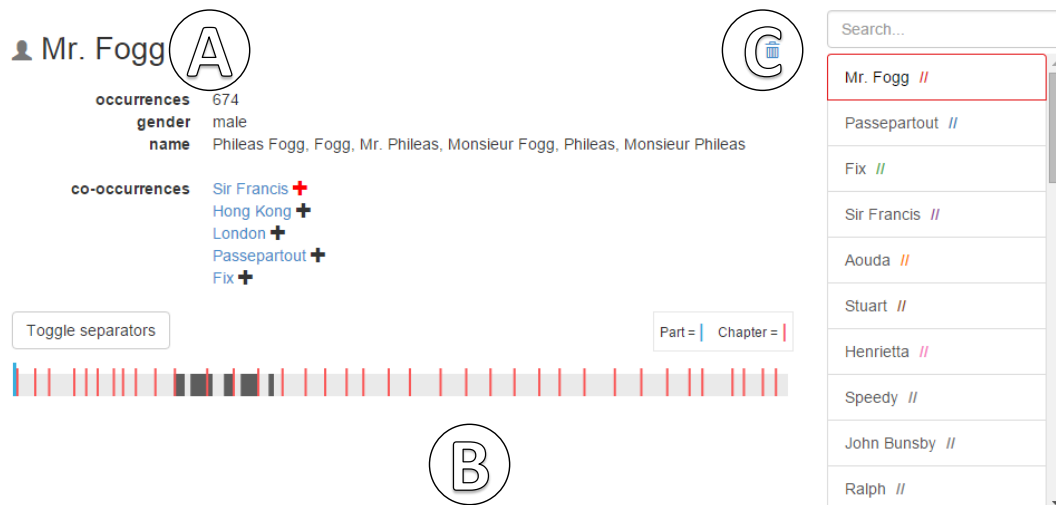


Figure 2: Character view listing information about the character *Mr. Fogg* extracted from the novel.

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 Figure 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 in at least one sentence, which is a common co-occurrence measure. However, other co-occurrence measures (number 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 represents connections between characters, as depicted in Figure 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 (Figure 4Ⓓ) to indicate where the characters occur in the novel, and with a range slider (Figure 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 (Figure 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 again available in that view (Figure 4Ⓐ), so that users can search, select and deselect the characters that are shown in the graph visualization. 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.



Figure 3: Combined fingerprint visualization showing two selected characters: *Passepartout* and *Fix*.

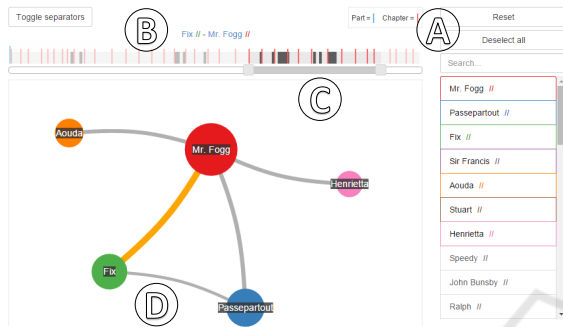


Figure 4: Character network visualization representing the co-occurrence graph of a selected text segment. This example represents connections between the main characters of the novel “The Lord of the Rings”. The nodes of the graph represent the characters and the edges the number of sentences in which they co-occur.

3.6 Word Cloud

If users are interested in what is being said near the characters or in getting a first impression of the contents of the novel, they can switch to the word cloud view. Word clouds are commonly used by literary scholars, as they are considered easy to understand (Clement et al., 2009) despite all their limitations (McNaught and Lam, 2010). The font size of the visualized words is scaled proportionally to their occurrence frequency, which is additionally shown in a tooltip 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 Figure 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 char-

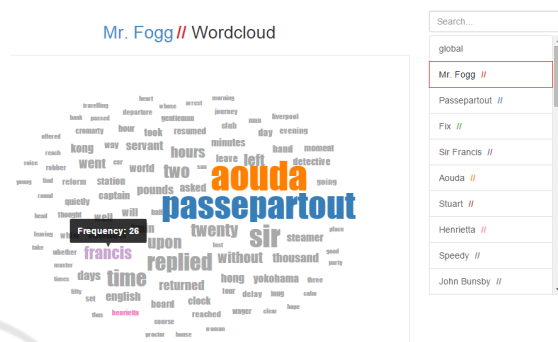


Figure 5: Word cloud view of the character *Mr. Fogg*.

acters. It reuses and extends an implementation of the University of Waterloo⁸, which takes annotation data and automatically generates narrative charts in the spirit of the aforementioned “Movie Narrative Charts” (Munroe, 2009).

Our approach adapts the visualization and displays the ten most frequently occurring characters as lines and every chapter as a node (cf. Figure 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 Figure 6, one can see, for example, that *Mr. Fogg* and *Fix* interact for the first time in chapter seven (Figure 6@), while *Aouda* and *John Bunsby* enter the plot in the middle of the novel (Figure 6@).

⁸<http://csclub.uwaterloo.ca/~n2iskand/?page.id=13/>




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

3.8 Text View

To support literary scholars in their common workflow, we also provide a text view where they can directly work with the text. Recognized chapters are listed as hyperlinks on the left (Figure 7(a)), while the text is presented on the right (Figure 7(b)). The focused chapter is emphasized with bold 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 a distant and close reading analysis. When hovering over the fingerprint blocks, the corresponding text passage is displayed in a tooltip (Figure 7(c)), and after clicking on one, the text view jumps to the corresponding position. Additionally, the literary scholars can jump 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 our approach. In the ePoetics project⁹, we developed approaches to support the analysis of German poetics – a form of early scholarly works on literature that

⁹<http://www.epoetics.de>

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 order 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 novels, since she read them some time ago, and is now trying to retrace the storyline and the important events by means of our approach.

4.1 Analysis of "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

Around the World in 80 Days

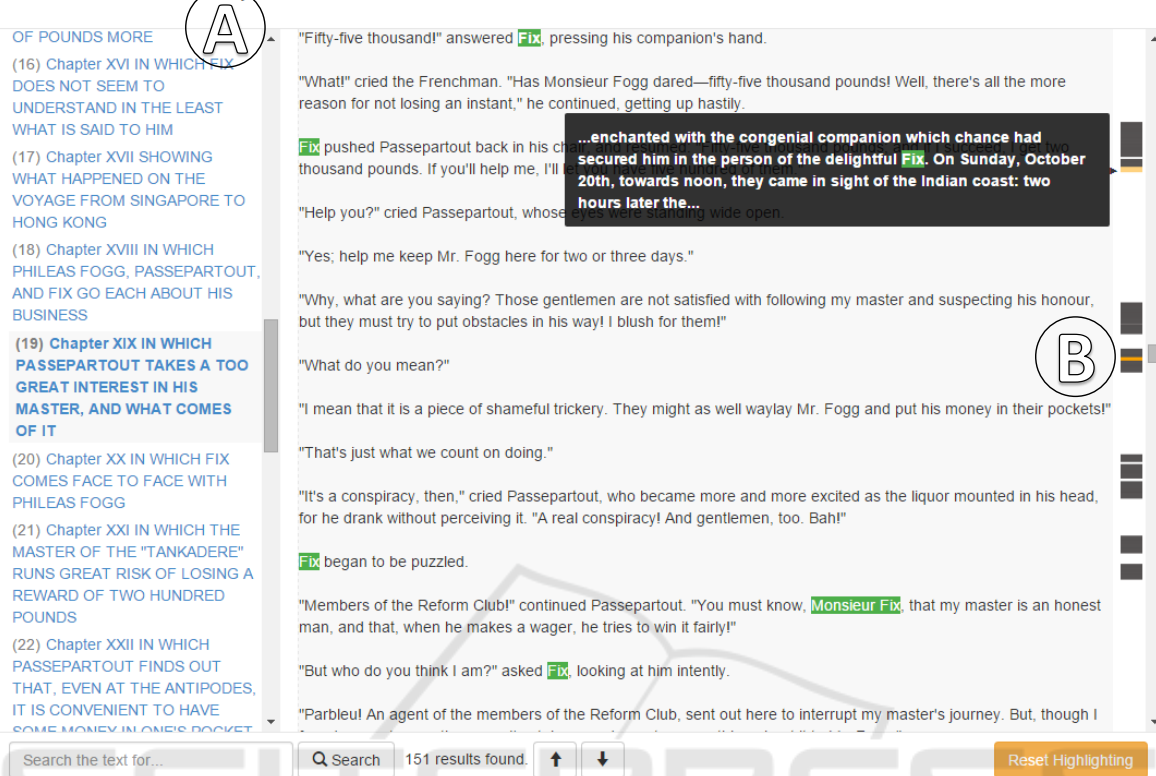


Figure 7: Text view with selected character *Fix*.

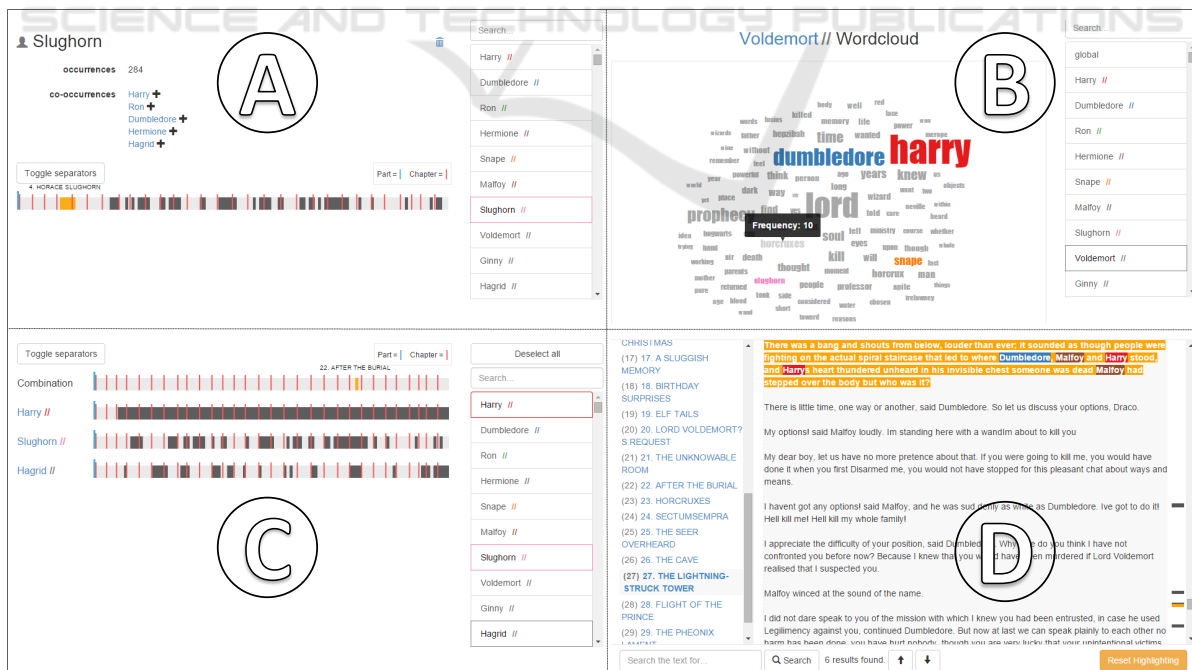


Figure 8: Some of the visualizations used by the literary scholar in the analysis of the novel “Harry Potter and the Half-Blood Prince”: character view (A), word cloud (B), fingerprint visualization (C) and text view (D).

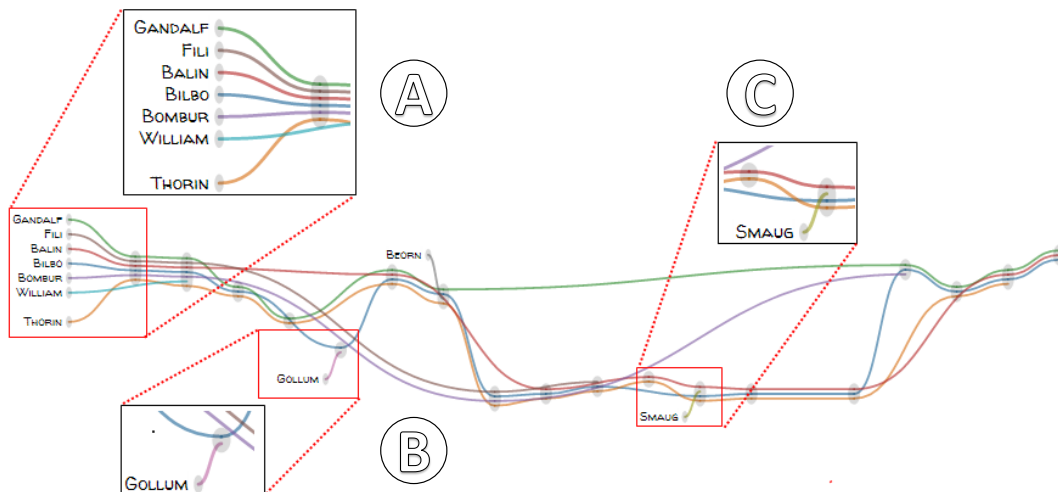


Figure 9: Plot view showing some of the main characters of J. R. R. Tolkien's novel "The Hobbit".

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 (Figure 8(a)) and uses the fingerprint visualization to jump directly to its first occurrence, which is opened in the 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 (Figure 8(b)).

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 (Figure 8(c)). 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 jumps directly to the text.

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 Figure 8(d). 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 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 Figure 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 (Figure 9(a)).

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 (Figure 9(b)) (origi-

nally 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 only a relationship between Gollum and Bilbo exist, as depicted in Figure 10(a). 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 (Figure 9(c)), 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 the 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 Thorin's, the legitimate king of Erebor's, chest within his tomb deep under the Lonely Mountain.

The usage scenarios show that our 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 confirming and generating hypotheses, confirming recollections and gaining insights.

5 DISCUSSION AND FUTURE WORK

Our approach was developed to support intra-document analysis by visualizing the dynamic relation between characters, and not for a comparative text of novels. However, we provide several visual abstractions, which offer different views on the text to support literary scholars in their analysis.

The approach can be flexibly extended by additional visual abstractions. For example, views that support literary scholars with visualizations of geographical information (Jänicke et al., 2012) could be integrated if the novel's setup is based on the real

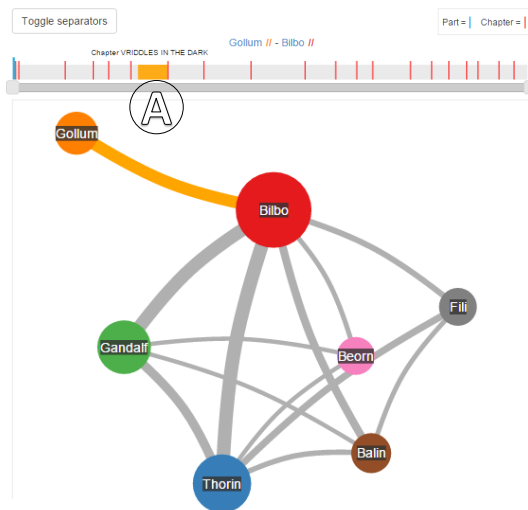


Figure 10: Character network view showing some of the main characters of “The Hobbit” and their co-occurrence frequencies.

world or offers own cartographic material.

In addition, the visual abstractions 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. This would enable literary scholars to recognize relationships between characters even better.

A further possibility might be to expand the plot view with automatically detected specific events (e.g., the “Battle of Helm’s Deep” in case of the novel “The Lord of the Rings”). Thereby, literary scholars could better track and identify the course of the story.

Our visual abstractions are based on out-of-the-box natural language processing toolkits and we are interested in examining how far we can get in supporting literary scholars with their analysis. Natural language processing methods have continuously improved and show strong robustness. Typically, they are rule-based, or rely on statistical and machine learning approaches, which have been trained on available modern newspaper or journal article texts, since large trainings sets are most often available for these types of texts. As a consequence, natural language processing techniques could be less effective when applied to a very specific type of text, such as historical or fictional novels.

This can lead to uncertainties and errors in the visual abstractions. As a consequence, uncritical interpretation of such visual representations can cause confusion and misunderstandings. Visualization can play an important role in addressing these problems

in the future. We intend to provide visual clues on the quality of automatic analyses (if available), for example, through color and opacity. This could help literary scholars to interpret automatically generated results. Furthermore, we want to let users steer and adapt techniques interactively according to their needs. As a result, automatic methods can be improved and applied to similar texts.

The current implementation of the detection of co-occurrences between entities is based on named entity recognition. We determine that two entities co-occur if they both appear at least in one sentence. This approach works well in many cases. However, it can be improved in various aspects. One option could be to let the literary scholars set the co-occurrences range variable, for example, several sentences or a whole paragraph. Additionally, methods for coreference resolution might be integrated, to detect more instances and compute more precise frequency values for entities. Coreference resolution is the task of resolving noun phrases to the entities that they refer to and there already exist robust methods (Raghunathan et al., 2010; Stoyanov et al., 2010). Existing approaches are usually trained on non-fictional texts and have to be adapted (where appropriate) to the users' needs.

In the future, we plan to take the semantics of the text into account in order to better investigate the relationships between entities. Again, already existing approaches (Gildea and Jurafsky, 2002; Ruiz-Casado et al., 2007) could be integrated into our approach to provide the literary scholars with more information for their analysis. In general, we aim at supporting a broader spectrum of natural language processing methods in the future.

Another issue is the challenge of scalability when working with arbitrarily long novels. Most of our visual abstractions scale very well. However, the fingerprint visualization and the plot view become less useful with longer text. Since the available screen space remains the same, they are increasingly compressed until they overlap. To address this problem in the future, we could integrate a focus+context technique, such as fisheye distortion (Bederson, 2000), to present detailed information in context, or an overview+detail approach to present multiple views with different levels of abstractions (Cockburn et al., 2009).

6 CONCLUSION

In this work, we presented an easily accessible web-based approach for visualizing the relation of characters and places in a novel. Our approach includes dif-

ferent natural language processing toolkits to extract named entities and possibilities to set parameters for variable analyses. Furthermore, we provide several visual abstractions, which support literary scholars with a flexible and comprehensive analysis of the novel characters. In particular, our approach facilitate distant reading and provide a starting point for new ideas, hypotheses and further analyses. All views enable literary scholars to jump to the corresponding text passage 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.

REFERENCES

- Abdul-Rahman, A., Lein, J., Coles, K., Maguire, E., Meyer, M., Wynne, M., Johnson, C. R., Trefethen, A., and Chen, M. (2013). Rule-based visual mappings – with a case study on poetry visualization. *Computer Graphics Forum*, 32(3pt4):381–390.
- Bederson, B. (2000). Fisheye menus. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*, pages 217–225. ACM.
- Clement, T., Plaisant, C., and Vuillemot, R. (2009). The story of one: Humanity scholarship with visualization and text analysis. In *Proceedings of the Digital Humanities Conference*, DH '09.
- Cockburn, A., Karlson, A., and Bederson, B. (2009). A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys*, 41(1):1–31.
- Don, A., Zheleva, E., Gregory, M., Tarkan, S., Auvil, L., Clement, T., Shneiderman, B., and Plaisant, C. (2007). 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 '07, pages 213–222. ACM.
- Dou, W., Wang, X., Skau, D., Ribarsky, W., and Zhou, M. (2012). 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 '12, pages 93–102.

- Eick, S., Steffen, J., and Sumner, E.E., J. (1992). Seesoft – a tool for visualizing line oriented software statistics. *IEEE Transactions on Software Engineering*, 18(11):957–968.
- Gildea, D. and Jurafsky, D. (2002). Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Hearst, M. A. (1995). Tilebars: Visualization of term distribution information in full text information access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '95, pages 59–66. ACM/Addison-Wesley.
- Heimerl, F., Lohmann, S., Lange, S., and Ertl, T. (2014). Word cloud explorer: Text analytics based on word clouds. In *47th Hawaii International Conference on System Sciences*, HICCS '14, pages 1833–1842. IEEE.
- Jänicke, S., Franzini, G., Cheema, M. F., and Scheuermann, G. (2015). On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges. In *Eurographics Conference on Visualization – STARs*, EuroVis '15. Eurographics Association.
- Jänicke, S., Geßner, A., Büchler, M., and Scheuermann, G. (2014). Visualizations for text re-use. In *Proceedings of the 5th International Conference on Information Visualization Theory and Applications*, IVAPP '14, pages 59–70. INSTICC.
- Jänicke, S., Heine, C., Stockmann, R., and Scheuermann, G. (2012). Comparative visualization of geospatial-temporal data. In *Proceedings of the 3rd International Conference on Information Visualization Theory and Applications*, IVAPP '14, pages 613–625. INSTICC.
- Keim, D. and Oelke, D. (2007). Literature fingerprinting: A new method for visual literary analysis. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, VAST '07, pages 115–122.
- Koch, S., Bosch, H., Giereth, M., and Ertl, T. (2011). Iterative integration of visual insights during scalable patent search and analysis. *IEEE Transactions on Visualization and Computer Graphics*, 17(5):557–569.
- Koch, S., John, M., Worner, M., Muller, A., and Ertl, T. (2014). Varifocalreader – in-depth visual analysis of large text documents. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1723–1732.
- Liu, S., Wu, Y., Wei, E., Liu, M., and Liu, Y. (2013). Storyflow: Tracking the evolution of stories. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2436–2445.
- McNaught, C. and Lam, P. (2010). Using wordle as a supplementary research tool. *The qualitative report*, 15(3):630–643.
- Moretti, F. (2005). *Graphs, maps, trees: abstract models for a literary history*. Verso.
- Munroe, R. (2009). Movie narrative charts. <http://xkcd.com/657/>.
- Oelke, D., Hao, M., Rohrdantz, C., Keim, D., Dayal, U., Haug, L., and Janetzko, H. (2009). Visual opinion analysis of customer feedback data. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, VAST '09, pages 187–194.
- Oelke, D., Kokkinakis, D., and Keim, D. A. (2013). Fingerprint matrices: Uncovering the dynamics of social networks in prose literature. *Computer Graphics Forum*, 32(3pt4):371–380.
- Phelan, J. (1989). *Reading people, reading plots: Character, progression, and the interpretation of narrative*. University of Chicago Press.
- Plaisant, C., Rose, J., Yu, B., Auvil, L., Kirschenbaum, M. G., Smith, M. N., Clement, T., and Lord, G. (2006). 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*, pages 141–150. ACM.
- Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., and Manning, C. (2010). A multi-pass sieve for coreference resolution. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, EMNLP '10, pages 492–501. ACL.
- Ruiz-Casado, M., Alfonso, E., and Castells, P. (2007). Automating the learning of lexical patterns: An application to the enrichment of wordnet by extracting semantic relationships from wikipedia. *Data & Knowledge Engineering*, 61(3):484–499.
- Stasko, J., Görg, C., and Liu, Z. (2008). Jigsaw: Supporting investigative analysis through interactive visualization. *Information Visualization*, 7(2):118–132.
- Stoyanov, V., Cardie, C., Gilbert, N., Riloff, E., Buttler, D., and Hysom, D. (2010). Coreference resolution with reconcile. In *Proceedings of the ACL 2010 Conference Short Papers*, pages 156–161. ACL.
- Tanahashi, Y. and Ma, K.-L. (2012). Design considerations for optimizing storyline visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2679–2688.
- Van Ham, F., Wattenberg, M., and Viegas, F. (2009). Mapping text with phrase nets. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1169–1176.
- Viégas, F. B. and Wattenberg, M. (2008). Tag clouds and the case for vernacular visualization. *interactions*, 15(4):49–52.
- Vuillemot, R., Clement, T., Plaisant, C., and Kumar, A. (2009). 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*, 2009, VAST '09, pages 107–114.
- Wattenberg, M. and Viegas, F. (2008). The word tree, an interactive visual concordance. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1221–1228.