

Added Value of Coreference Annotation for Character Analysis in Narratives

Melanie Andresen, Michael Vauth

Universität Hamburg, Technische Universität Hamburg
Hamburg, Germany

melanie.andresen@uni-hamburg.de, michael.vauth@tuhh.de

Abstract

A central question for the analysis of literary texts in digital humanities is the identification of characters in the text. A simple approach would be to reduce the analysis to mentions of the character by its proper name as these are easy to retrieve. A more elaborate solution requires coreference annotation, identifying all mentions of a character in the text, irrespective of their form. Using the example of the novel *Corpus Delicti* by the German author Juli Zeh, we compare these two approaches and show the added value of coreference annotation.

Keywords: coreference, character analysis, literary studies, narrative texts

1. Introduction

One of the objectives of the analysis of literature in digital humanities (DH) is the analysis of characters (Jannidis, 2009). This analysis can mainly focus on two aspects: First, we can be interested in presence and copresence of characters in the course of the text. Whenever a character is mentioned in the text, we say it is present in this part of the text. Whenever two or more characters are mentioned in a defined window of k words, sentences or paragraphs, we say that they are copresent. Second, we can be interested in characterization and want to learn about properties of a character. In this case, we can investigate what is said about the character in the text.

Both types of analysis require the detection of character mentions in the text. One simple solution would be to reduce the analysis to mentions of the proper name of the character. These can be identified easily and can be considered an approximation for the actual presence of a character. The more time-consuming way is a coreference annotation of the text. Two expressions in a text are considered coreferent if they refer to the same discourse entity (e. g. Kübler and Zinsmeister, 2015). In addition to proper names, this includes pronouns and noun phrases.

In this paper, we compare the approaches with and without coreference annotation in order to show the added value of this annotation. The novel *Corpus Delicti* by the German author Juli Zeh is the example text basis of our analysis. We will show that proper names are only a small part of character mentions. Moreover, the distribution of proper names vs. pronouns varies in the text and some types of text such as conversation are underrepresented when focusing on proper names only. Our focus is on character presence, characterization is discussed briefly.

2. Related Work

2.1. Coreference Annotation

Despite extensive research on the automation of coreference resolution, the evaluation scores for the best results still range between 70 and 80 (MUC and B^3 , Lee et al., 2018). The automatic detection of character mentions in literary texts is considered to be especially challenging, because references by forms other than named entities are frequent (Vala et al., 2015).

Vala et al. (2015) propose a system for character identification in novels. However, they only aim at extracting a list of all characters and names used to refer to these. The results range between F1-scores of 0.45 and 0.76. They use the resulting lists to compare the number of characters in novels from a diachronic perspective and in novels with an urban vs. a rural setting. In both cases, they find no significant differences. Given the low scores of the automation task, Vala et al. (2016) focus on manual annotation and present an annotation tool for this purpose.

For German, there is the statistical coreference resolution tool HotCorefDE (Rösiger and Kuhn, 2016) and a rule-based approach tailored to historic literary text (Krug et al., 2015), which achieve competitive results. However, in order to make our comparison of analyses based on proper names vs. coreference annotation meaningful, we need an annotation quality that can currently only be achieved by manual annotation.

2.2. Character Networks

Our work is situated in the context of character network analysis, made popular by e. g. Moretti (2011). Piper et al. (2017) present recent approaches to investigating the historical development of character networks using graph metrics and revisit the notion of interaction by asking readers of fiction to provide interaction labels. Xanthos et al. (2016) present an approach to include the development of character networks over time in the text.

For German, several studies on character networks have been conducted and some tools have been published. The rCat-Tool¹ by Barth et al. (2018) visualizes a character network taking a narrative text and a list of character names as input. Additionally, it creates word-clouds of the most frequent words appearing near to the mentions of a specific character. It can take several names for one character into account, but does not resolve pronouns or noun phrases not explicitly provided.

Blessing et al. (2017) create a character network for the Middle High German *Parzival* only taking named entities and noun phrases into account. Similar to our approach,

¹<http://www.ims.uni-stuttgart.de/forschung/ressourcen/werkzeuge/rcat.html>, May 2, 2018.

they compare an analysis of proper names to an analysis of both proper names and noun phrases and inspect the influence of direct speech and embedded entities (e. g. in possessive constructions). Krautter (2018) analyses copresence in a dramatic text and bases his analysis on character speech explicitly attributed to a speaker, which is hardly possible for narrative texts.

The research project Kallimachos focuses on the analysis of character networks in narrative texts, namely German novels (Jannidis et al., 2015; Jannidis et al., 2016). They also exploit the benefits of coreference annotation and their preprocessing pipeline includes a tool for automatic coreference resolution (Krug et al., 2015). They provide a preprocessing pipeline for tagging, parsing, named-entity resolution and coreference resolution and a python tutorial² for generating character networks from the result.

3. Data and Data Annotation

Our object of analysis is Juli Zeh’s novel *Corpus Delicti* that was published in 2009. The story is set in a dystopian future where a repressive system is in power that is centered around questions of human health. The novel is written in a realistic style and narrated by a heterodiegetic narrator who rarely imparts psychological information about the characters. Many chapters include large dialog sections (about half of the novel is direct speech) and every chapter represents a fairly self-contained story episode, which is why we base our analyses on a segmentation per chapter.

We conducted the manual annotation using the tool CorefAnnotator³ by Nils Reiter that was developed specifically for the task of manual coreference annotation. For the annotation we used the guidelines for coreference annotation described in Rösiger et al. (2018). In contrast to these guidelines, the annotation task was restricted to the annotation of characters, i. e. mentions of humans. For this reason, we expected the annotation task to be rather unambiguous and did rely on single annotation. The text (about 46.000 token) was split between four annotators. Ambiguous instances were marked by the individual annotators during the annotation process and discussed in the group. In the end, one of the annotators merged the four sections and merged the mention sets of characters that appear in more than one section.

In addition, the text was annotated automatically for part-of-speech with MarMot (Müller et al., 2013), lemma and dependency syntax using MATE (Bohnet, 2010) with a model trained on the *Hamburg Dependency Treebank* (Foth et al., 2014). We have previously evaluated the quality of pos tagging and dependency parsing for literary data (Adelmann et al., 2018a; Adelmann et al., 2018b) and these tools emerged as the most reliable. On a text section of the novel *Corpus Delicti* (1,518 token), MarMot achieved an accuracy of 0.97, and MATE a labeled attachment score of 0.87 (unlabeled: 0.91).

From this annotated version of the novel, we extracted the following information for each character mention:

- the token span,
- the entity it refers to,
- the linguistic form (proper name, pronoun...),
- whether it occurs inside direct speech (detected by quotes) and
- the chapter in which it occurs.

For copyright reasons, we are unable to publish the annotated text. However, you can find the extracted list of mentions and their annotations at <https://doi.org/10.5281/zenodo.1239701>.

4. Results

4.1. Form of Reference

We will first inspect the forms used to refer to the characters in the novel. Figure 1 displays the distribution of linguistic forms used to refer to the four characters mentioned most frequently in the novel (the white line indicates the absolute numbers, scale on the right). The distribution of forms is very similar for the four characters. Proper names (NE) amount to about one quarter of all mentions, while personal pronouns (PPER) account for about half of all mentions. The rest are mostly possessive pronouns (PPOSAT) and noun phrases (NP). We conclude that while the relation between proper names and other mentions is relatively stable across characters in our data, proper names account for only a minor proportion of all character mentions, making a coreference annotation desirable.

In order to see if the relation between proper names and other mentions is also stable across the text, we display the distribution of expressions referring to the main character of *Corpus Delicti*, Mia, across chapters in Figure 2. In the horizontal dimension, we can see the 49 chapters of the novel. For each of the chapters the bars display which forms are used for reference to Mia, relative to all references to Mia in the chapter (scale on the left). The white line graph indicates the absolute number of mentions of Mia in each chapter (scale on the right).




Figure 1: Forms used to refer to main characters of *Corpus Delicti*: x-axis: the four main characters, y-axis, bars: proportion of the forms (left scale), y-axis, line: absolute number of mentions per character (right scale).

²http://kallimachos.de/kallimachos/index.php/Tutorial_Figurennetzwerke, May 2, 2018.

³<https://doi.org/10.5281/zenodo.1228105>

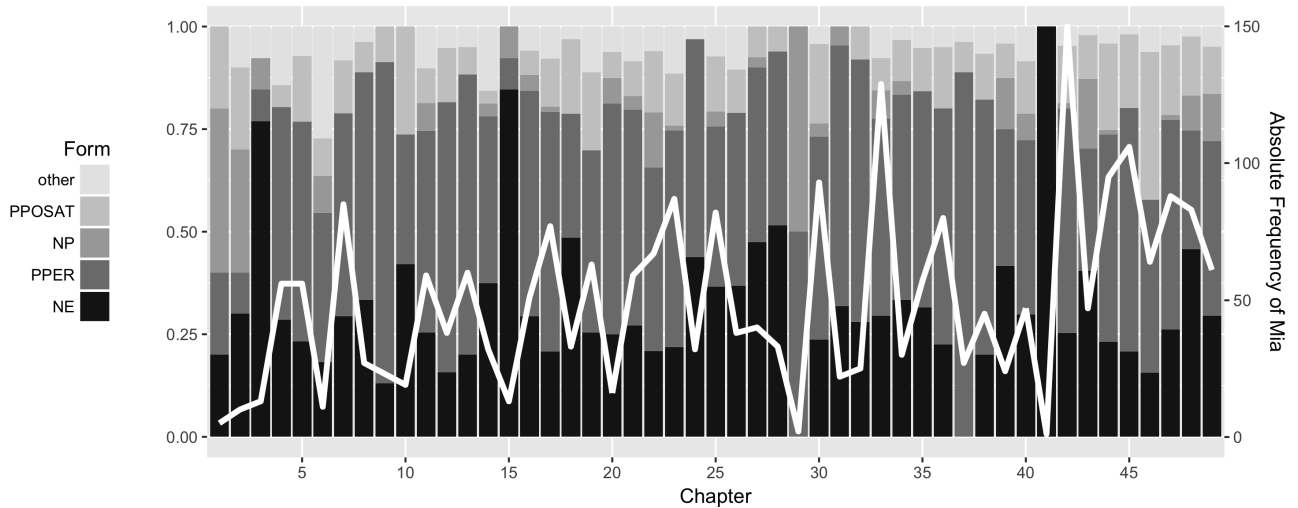


Figure 2: Forms used to refer to the main character of *Corpus Delicti*, Mia, by chapters: x-axis: the 49 chapters of the novel, y-axis, bars: proportion of the forms (left scale), y-axis, line: absolute number of mentions of Mia (right scale).

The distribution of forms for the character Mia shows a considerable amount of variability. Note that the main outliers occur in chapters with very few mentions of the character. We can posit the hypothesis that the relation between different forms of referring expressions is increasingly stable when including longer stretches of text. It can also be noted (having read the novel) that the proportion of proper names is high in chapters with other female characters which make pronouns ambiguous. This applies to, for example, chapters 3 and 28.

4.2. Character Distributions with and without Coreference Annotation

For studies on the presence of characters, e. g. with the purpose of creating a character network, we need information on when characters appear in the novel and which other characters they appear with. We want to describe how this information is affected by preprocessing such as coreference annotation. In addition to this main question, we also explore the effect of excluding direct speech (see also Blessing et al., 2017).

In Figure 4 you can see a comparison of three conditions, ranging from no preprocessing (proper names only, Figure 4a) via one preprocessing step (with coreference annotation, Figure 4b) to two preprocessing steps (coreference annotation and exclusion of direct speech, Figure 4c). The bars indicate the relative proportion of character mentions for each of the four main characters, relative to all mentions of all four characters in the chapter.

One expected, but still important change between the three conditions is the number of mentions (white line, scale on the right). When including coreference annotation, the number of mentions increases greatly. This is highly beneficial for analysis of, for example, syntactic contexts the characters occur in. Naturally, the number of mentions decreases again when excluding direct speech.

When comparing the first two graphs, we can see some changes in the proportions of the four characters. For instance, the relation between the characters Mia and Moritz

is almost inverted in chapter 27. A closer look at the chapter reveals the reason: A large part of the chapter is a conversation between Mia and Moritz about something that happened to Moritz. In the conversation, references to Moritz are mainly realized by first and second person pronouns. Another example is chapter 3: As you can see, the proportion of the mentions of Kramer rises relative to the mentions of Mia when coreference annotation is included. The reason is that Kramer is having a conversation in the chapter, so he is often addressed by pronominal expressions in first and second person. Mia, on the other hand, is not communicatively interacting with the other characters but the conversation’s topic. For that reason, Mia nearly ‘disappears’ when direct speech is excluded (see c).

We can derive the hypothesis that references to speaker and addressee in character conversation are rarely realized as proper names and will often be underrepresented if only proper names are considered. In addition to that, the presence of an absent character like Mia would be overestimated.

The distribution changes even more when all mentions in direct speech are excluded (Figure 4c). Generally speaking, the number of characters per chapter is reduced. Most strikingly, the proportions of the character Moritz are considerably smaller. This can be explained by his position in the novel: The character Moritz died before the main narration and is – except for some flashbacks – not present but only talked about.

It is a matter of the individual focus of analysis whether the exclusion of direct speech is appropriate or not, e. g. the fact that Mia talks a lot about her dead brother Moritz does also tell us a great deal about the character constellation. However, we have to bear in mind that this decision can heavily influence the analysis.

In addition to the visual comparison of the graphs, Figure 3 provides Pearson’s correlations for the absolute frequency of mentions of the four characters under the three conditions (in the same order as in Figure 4). A score of 1 would mean that there is a linear relationship between the

	Mia	Kramer	Rosentreter	Moritz
proper names only				
coreference	0.87	0.94	0.94	0.90
coreference, no direct speech	0.84	0.93	0.97	0.89

Figure 3: Pearson’s correlations for the absolute frequency of mentions per chapter for the four characters under the three conditions

frequencies: If the frequency of a character measured in proper names doubles from one chapter to another, the corresponding frequency measured in total mentions including pronouns etc. doubles as well. The further away from 1 the score is, the more independent the frequencies are. If the correlation between two conditions is strong, we do not get much additional information when considering both conditions.

We can see that the characters are affected by the changes to different degrees. For Kramer and Rosentreter all the scores are above 0.93, indicating some, but no substantial changes. The scores for Mia and Moritz are much lower (between 0.78 and 0.90). The comparison of the first two conditions reflects the variability we have already seen for Mia in Figure 2: The relation of proper names to other types of mentions is not constant across chapters.

Also the correlations between condition 2 (coreference) and condition 3 (coreference, direct speech excluded) are lower for Mia and Moritz. This can be explained by the fact that they are more frequently mentioned in direct speech than Kramer and Rosentreter.

The correlation between condition 1 (proper names only) and condition 3 (coreference and direct speech excluded) is especially low for Mia and Moritz. This indicates that the two preprocessing steps—coreference resolution and exclusion of direct speech—accumulate for these characters.

4.3. Characterization by Noun Phrases

In addition to increasing the accuracy of analyses based on character mentions, coreference annotation can also give us a first impression of a character’s attributes. For this purpose, we inspect the noun phrases used to refer to a specific character. Note that we use the term ‘noun phrase’ in a narrow sense, referring to appellative noun phrases as opposed to proper names and pronouns. In Table 1 you can see the most frequent noun phrases used for the character Mia, reduced to the head of the phrase. We can see immediately that the story of Mia is centered around court proceedings in which Mia is the defendant and – in the end – the convicted. Additionally, we can see that her family status as a sister is mentioned repeatedly. This is mirrored in the data for her brother Moritz: 43 of 47 noun phrases referring to him have the head *Bruder* (‘brother’). Here we can see clearly that while the main character Mia is characterized by many different noun phrases, Moritz’s role is limited to his family relation to Mia.

Noun Phrase	Translation	Frequency
Angeklagte	defendant	32
Schwester	sister	7
Beschuldigte	accused	7
Verurteilte	convicted	6
Mandantin	client	4

Table 1: The most frequent appellative noun phrases used to refer to the main character Mia

Other possibilities for the automatic description of a character include word clouds based on the context of all character mentions as in Barth et al. (2018). Our future focus, however, will be on more linguistically informed approaches that rely on syntactic annotations. In this way, we can extract explicit attributions as realized in non-verbal predicates (e. g. *Rosentreter ist ein guter Junge*, ‘Rosentreter is a good boy’) or all full verbs used with a selected character as subject.


5. Conclusions

We have shown that the ratio of proper names to other mentions is about 1:3. While it is surprisingly stable between characters, it varies between chapters of the novel *Corpus Delicti*. We could show that especially speakers in conversation are highly underrepresented when considering proper names only. For this reason, the analysis of copresence of characters will yield different results when based on proper names only or on all mentions. We therefore argue that, while the analysis of proper names requires much less time, literary character analysis benefits from coreference annotation. In addition, it enhances the possibilities of describing a character by the noun phrases referring to it and its syntactic context.


In the future, we will use the data described here to create character networks and further investigate the influence of preprocessing like coreference annotation for this type of analysis. We hope that this type of analysis will contribute to the identification of genre features, our focus being the dystopia. To allow for genre specific findings, we will annotate and analyze coreference in another dystopia as well as two historic novels.

6. Acknowledgements


This work has been funded by the ‘Landesforschungsförderung Hamburg’ in the context of the *hermA* project (LFF-FV 35). We thank Lea Röseler and Daniel Fabian Klein for their help with the annotation and Piklu Gupta for checking our English. All remaining errors are our own.



(a) based on proper names only



(b) with coreference annotation



(c) with coreference annotation, direct speech excluded

Figure 4: Mentions of main characters by chapters under different conditions

7. Bibliographical References

- Adelmann, B., Andresen, M., Menzel, W., and Zinsmeister, H. (2018a). Evaluating Part-of-Speech and Morphological Tagging for Humanities' Interpretation. In *Proceedings of the Second Workshop on Corpus-Based Research in the Humanities*, pages 5–14, Vienna, Austria.
- Adelmann, B., Andresen, M., Menzel, W., and Zinsmeister, H. (2018b). Evaluation of Out-of Domain Dependency Parsing for its Application in a Digital Humanities Project. In *Proceedings of the 14th Conference on Natural Language Processing (KONVENS 2018)*, Vienna, Austria.
- Barth, F., Kim, E., Murr, S., and Klinger, R. (2018). A reporting tool for relational visualization and analysis of character mentions in literature. In *Book of Abstracts of DHD 2018*, pages 123–127, Cologne, Germany.
- Blessing, A., Echelmeyer, N., John, M., and Reiter, N. (2017). An End-to-end Environment for Research Question-Driven Entity Extraction and Network Analysis. In *Proceedings of the Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 57–67, Vancouver, Canada.
- Bohnet, B. (2010). Very High Accuracy and Fast Dependency Parsing is not a Contradiction. In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING 2010)*, Beijing, China.
- Foth, K., Köhn, A., Beuck, N., and Menzel, W. (2014). Because Size Does Matter: The Hamburg Dependency Treebank. In *Proceedings of LREC 2014*, pages 2326–2333, Reykjavik, Iceland.
- Jannidis, F., Reger, I., Weimer, L., Krug, M., Toepfer, M., and Puppe, F. (2015). Automatische Erkennung von Figuren in deutschsprachigen Romanen. In *Book of Abstracts of DHD 2015*, Graz, Austria.
- Jannidis, F., Reger, I., Krug, M., Weimer, L., Macharowsky, L., and Puppe, F. (2016). Comparison of Methods for the Identification of Main Characters in German Novels. In *Digital Humanities 2016: Conference Abstracts*, pages 578–582, Kraków.
- Jannidis, F. (2009). Character. In Peter Hühn, et al., editors, *Handbook of Narratology*, number 19 in Narratology: contributions to narrative theory, pages 14–29. de Gruyter, Berlin [u.a.].
- Krautter, B. (2018). Quantitatives close reading? Vier mikroanalytische Methoden der digitalen Dramenanalyse im Vergleich. In *Book of Abstracts of DHD*, pages 295–300, Cologne, Germany.
- Krug, M., Puppe, F., Jannidis, F., Macharowsky, L., Reger, I., and Weimar, L. (2015). Rule-based Coreference Resolution in German Historic Novels. In *Proceedings of the Fourth Workshop on Computational Linguistics for Literature*, pages 98–104.
- Kübler, S. and Zinsmeister, H. (2015). *Corpus Linguistics and Linguistically Annotated Corpora*. Bloomsbury, London, New York.
- Lee, K., He, L., and Zettlemoyer, L. (2018). Higher-order Coreference Resolution with Coarse-to-fine Inference. Accepted at NAACL 2018. *arXiv:1804.05392*.
- Moretti, F. (2011). Network Theory, Plot Analysis. *New Left Review*, 68:80–102.
- Müller, T., Schmid, H., and Schütze, H. (2013). Efficient Higher-Order CRFs for Morphological Tagging. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 322–332, Seattle, Washington, USA, October. Association for Computational Linguistics.
- Piper, A., Algee-Hewitt, M., Sinha, K., Ruths, D., and Vala, H. (2017). Studying Literary Characters and Character Networks. In *Digital Humanities 2017, Conference Abstracts*, pages 119–122, Montreal, Canada.
- Rösiger, I. and Kuhn, J. (2016). IMS HotCoref DE: A Data-driven Co-reference Resolver for German. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*.
- Rösiger, I., Schulz, S., and Reiter, N. (2018). Towards Coreference for Literary Text: Analyzing Domain-Specific Phenomena. In *Proceedings of LaTeCH-CLfL*.
- Vala, H., Jurgens, D., Piper, A., and Ruths, D. (2015). Mr. bennet, his coachman, and the archbishop walk into a bar but only one of them gets recognized: On the difficulty of detecting characters in literary texts. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 769–774.
- Vala, H., Dimitrov, S., Jurgens, D., Piper, A., and Ruths, D. (2016). Annotating characters in literary corpora: A scheme, the charles tool, and an annotated novel. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France.
- Xanthos, A., Pante, I., Rochat, Y., and Grandjean, M. (2016). Visualising the dynamics of character networks. In *Digital Humanities 2016: Conference Abstracts*, pages 417–419, Krakow.