

# Automated Event Annotation in Literary Texts

Michael Vauth<sup>1</sup>, Hans Ole Hatzel<sup>2</sup>, Evelyn Gius<sup>1</sup> and Chris Biemann<sup>2</sup>

<sup>1</sup>*Technical University of Darmstadt, Institute of Linguistics and Literary Studies, Dolivostraße 15, 64293 Darmstadt, Germany*

<sup>2</sup>*Universität Hamburg, Language Technology Group, Vogt-Kölln-Straße 30, 22527 Hamburg, Germany*

## Abstract

We approach the modeling of event structure of literary texts with narratological event concepts. A manually annotated corpus of 4 prose texts with event categories allows us to learn to automatically classify events using a transformer model, relying on a rule-based system in conjunction with a pre-trained parser to identify events. For the evaluation of both manual and automated annotation, we use narrativity graphs, which capture the change in narrativity over the course of the text. In an exploratory analysis, we apply the event classifier in conjunction with graph-based narrativity metrics to a large literary corpus. We find that text length does neither influence the length of eventful passages nor the number of eventful passages in the beginnings of texts.

## Keywords

event, narrativity, automation, annotation, literary studies, narrative theory

## 1. Introduction

We present an approach to the automation of the annotation of narratological event concepts. By using concepts of literary theory for the annotation of events, we aim at an added-value analysis for the interests of literary studies, where events are conceived as minimal narrative units. An automated approach enables a large-scale analysis of event-related patterns in literary texts. We first introduce theoretical concepts of events in narrative theory and the approach to events in natural language processing (Section 2). Then, we describe our scheme for the manual annotation of events together with possible evaluations of these annotations (Section 3). In Section 4, we present our transformer-based method for automatic event annotation. Finally, we present some exploratory corpus analyses based on the automatic event annotations to show its application potential (Section 5). In this work, we lay the foundations for narrativity-based computational event representations. Shifting the focus away from lexical semantics (cf. [9, 3]) towards the narrative impact of events allows us to set aside modeling the participants and arguments of, and the temporal relations between events.

---

*CHR 2021: Computational Humanities Research Conference, November 17–19, 2021, Amsterdam, The Netherlands*

✉ michael.vauth@tu-darmstadt.de (M. Vauth); hans.ole.hatzel@uni-hamburg.de (H.O. Hatzel); evelyn.gius@tu-darmstadt.de (E. Gius); christian.biemann@uni-hamburg.de (C. Biemann)

🆔 0000-0002-3668-6273 (M. Vauth); 0000-0002-4586-7260 (H.O. Hatzel); 0000-0001-8888-8419 (E. Gius); 0000-0002-8449-9624 (C. Biemann)

© 2021 Copyright for this paper by its authors.  
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

**Table 1**  
Event Types in Narratology.

|                                     | Event Types   |  |
|-------------------------------------|---|--|
|                                     | Less Narrative/Eventful   | More Narrative/Eventful  |
| Schmid [24]                         | <b>Change of state:</b> “(1) a temporal structure with at least two states, the initial situation and the final situation; and (2) the equivalence of the initial and final situations, that is, the presence of a similarity and a contrast between the states, or, more precisely, the identity and difference of the properties of those states” | <b>Event:</b> exceptional (relevant, unpredictable, persistent, irreversible, non-iterative) change of state                                 |
| Lahn/Meister [16]<br>Ryan [23]      | <b>Happening:</b> expected change of state<br><b>Happening:</b> “occur accidentally, having a patient but not an animated agent”  | <b>Event:</b> unexpected change of state<br><b>Action:</b> “targeted toward a goal and have a voluntary human or human-like agent”           |
| Martinez/Scheffel [18]<br>Ryan [22] | <b>Happening:</b> unintended change of state<br><b>Change of physical states:</b> “deliberate actions and accidental happenings”  | <b>Action:</b> intended change of state<br><b>Mental act:</b> “mental acts can be regarded as a hybrid of transient event and durable state” |
| Prince [21]                         | <b>Stative event:</b> “any event which describes a state”   | <b>Active event:</b> “any event which describes an action”, where action means any process in time   |

## 2. Event Concepts in Narratology and NLP

### 2.1. Event Concepts in Narratology

Most concepts of narrative theory are based on the distinction between story (*histoire*) and textual representation (*discourse*) [4]. For that reason, events, as minimal units of the story, can be seen as foundational categories of narratology [14]. Nevertheless, apart from Meister [19], there have been no efforts to operationalize these event concepts for annotations. This is probably due to events often being defined not as an analytical but rather as a heuristic category [10, p. 233]. In earlier contributions to digital narratology, events and associated narrativity have therefore been assessed as difficult categories to operationalize [10, p. 244].

Table 1 shows common conceptions of event in narrative theory. Most of these are based on a notion of change of state and introduce additional parameters to differentiate event types. For example, Schmid [24] introduces five event properties to evaluate events with regard to their narrativity. Following this proposal, a change of state’s narrativity rises with its relevance for the narrated story, with its unpredictability in regard to the past course of the story, with its persistence for the following story, with its irreversibility or with its singularity. Other concepts take into account whether a change of state is the result of an intentional act by an anthropomorphic agent. With Prince [21] and his concepts of stative and active events, there is only one theorist who does not define events as a change of states. This grounds on a textual definition of events: Prince has suggested, “to call event in a story any part of that story which can be expressed by a sentence, where sentence is taken to be the transform of at least one, but less than two, discrete elementary strings” [21, p. 17].

## 2.2. Event Detection in NLP

In natural language processing (NLP), events are typically understood as anything that is “happening during a particular interval of time” [12], which relates closely to Prince’s broad definition of active events (see Table 1). Originally, NLP approaches to event detection focused on the news domain [5], with more recent work adapting the concept to other domains such as biomedicine, information security and literature [2, 17, 25]. In the popular ACE approach [8], events were annotated with regard to a specific set of 28 semantic categories in the news domain, enabling a supervised machine learning approach to modeling event semantics with high coverage [27].

In the literature domain, Sims et al. [25] differentiate between realis and irrealis events, thereby distinguishing between events actually happening in the narrated world and those that, for example, only happen in the thoughts of characters. They make use of event detection to establish a correlation between the frequency of realis events in a text and its reception. Recognizing the limitations of fixed event categories in the more open domain of literature, they adopt an open approach to event categories.

However, while the distinction of events and non-events is typically made, the impact of event types, or categories, on the narrative is not explicitly modeled. We limit the event trigger detection to verbs rather than including nouns and adjectives.

## 3. Manual Event Annotation in Digital Narratology

### 3.1. Our Annotation Schema

While most definitions of events in narrative theory rely on change of states (see Table 1), we adopted a more fine grained approach. Our classification of events is based on core features of events in narrative theory (i.e., being a state, a process in time and change of state). This also allows us to include processes of speaking, thinking and movement.<sup>1</sup>

Additionally, we adopted Prince’s [20] idea that events can be presented by a single sentence by analyzing verbal phrases.

**Annotation Spans.** We annotate each verbal phrase. In cases where a text segment cannot be assigned to a verbal phrase, it is annotated as <non\_event> (see below). To avoid overlapping annotations, an annotation span is defined by an inflected verb and its direct subordinated dependencies. Subordinated verbal phrases do not lead to overlapping annotations:

1. [1As Gregor Samsa one morning from uneasy dreams awoke]1, [2found he himself in his bed into a monstrous insect-like creature transformed]2.<sup>2</sup>

**Event Types.** Our annotation scheme provides four event types: changes of state, process events, stative events and non-events. A given verbal phrase’s event type is defined by the act inside the narrated world that the full verb signifies.

---

<sup>1</sup>These do not imply a change of state as defined in narratological event concepts, however they are part of Chatman’s [4] enumeration of principal components of action: “The principal kinds of actions that a character or other existent can perform are nonverbal physical acts [...], speeches [...], thoughts [...], and feelings, perceptions, and sensations.”[4, p. 45]

<sup>2</sup>This and all following examples are taken from Franz Kafka’s narration *Die Verwandlung* (Translations taken from Kafka [15]).

**Table 2**

Event type distribution in the manually annotated corpus

|                 | <i>Die Verwandlung</i> |         | <i>Das Erdbeben</i> |       | <i>Effi Briest</i> |         | <i>Krambambuli</i> |       |
|-----------------|------------------------|---------|---------------------|-------|--------------------|---------|--------------------|-------|
| non_event       | 0.28                   | (653)   | 0.24                | (175) | 0.42               | (2,843) | 0.32               | (187) |
| stative_event   | 0.19                   | (444)   | 0.17                | (120) | 0.27               | (1,838) | 0.18               | (103) |
| process_event   | 0.52                   | (1,221) | 0.57                | (406) | 0.30               | (2,017) | 0.49               | (282) |
| change_of_state | 0.01                   | (34)    | 0.01                | (9)   | 0.01               | (47)    | 0.01               | (3)   |

The event type <**change\_of\_state**> covers physical and mental state changes of animate and inanimate entities. The change of state needs to be expressed in a single verbal phrase. An example of a (mental) change of state is the process of waking up, described by the first verbal phrase in Example 1.

The event type <**process\_event**> covers actions and happenings that do not lead to a change of state, such as processes of moving, talking, thinking and feeling. For example, the second verbal phrase in Example 1 is a process of perception (“found [...] himself [...] transformed”) and, therefore, has to be annotated as a process event.<sup>3</sup>

The event type <**stative\_event**> covers all verbal phrases that refer to the state of an animate or inanimate entity. This also includes physical and mental states. Unlike changes of state and processes, stative events do never refer to temporal processes. As Example 2 shows, stative events are often descriptions of places, persons or objects.

2. [<sub>1</sub>There stood a bowl filled with sweet milk]<sub>1</sub> [<sub>2</sub>in which swam small bits of white bread.]<sub>2</sub>

The event type <**non\_event**> finally covers verbal phrases that do not refer to a fact in the narrated world. The most common variants of non-events are questions or modalized and generic statements (Example 3). Non-events are a necessary addition allowing us to annotate all text passages not just those referring to events, thereby simplifying the annotation schema.

3. [<sub>1</sub>A man must have his sleep.]<sub>1</sub>

### 3.2. Corpus, Manual Annotations and Inter Annotator Agreement

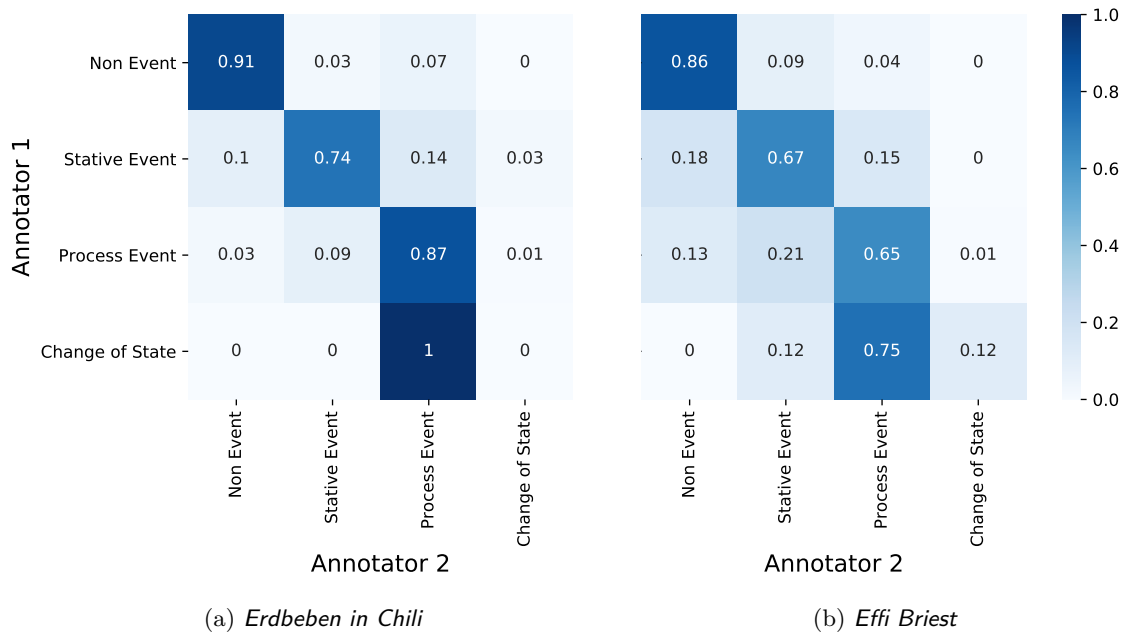
So far, we have double annotated four German prose texts with 151,000 tokens in total: Besides Kafka’s *Metamorphosis*, these are Theodor Fontane’s *Effi Briest* (partially; 14 of 36 chapters), Marie von Ebner-Eschenbach’s *Krambambuli* and Heinrich von Kleist’s *Das Erdbeben in Chili*.<sup>4</sup> After some training, the annotators (students in literary studies) reached an agreement of 0.73 Krippendorff’s  $\alpha$ .<sup>5</sup> Figure 1 shows a confusion matrix for two of the texts. The agreement differs both depending on the categories / event types and the annotated text. We suspect that the agreement is influenced by the linguistic style of the respective authors.

Our manual annotations show that the first three event types are the most common: non-events, stative events and process events are the most common types as Table 2 shows. For the three shorter texts, *Die Verwandlung*, *Das Erdbeben in Chili* and *Krambambuli*, even the

<sup>3</sup>This example shows why it is important to use the full verb as an annotation criterion. The transformation is only described as an object of perception and not as an objective fact in the narrated world.

<sup>4</sup>These annotations are quite time-consuming. For the four sections or the 14 chapters of *Effi Briest*, for example, each of our annotators needed over 60 working hours.

<sup>5</sup>The annotation’s span overlap average is 97%.



**Figure 1:** Confusion matrices for manual annotations in *Erdbeben in Chili* and *Effi Briest* with row normalized values. Compare Figure 4 for the automated annotations.

proportion between the first three types is similar. Only in Theodor Fontane’s novel *Effi Briest* non-events are the most common event type. This is due to the high proportion (here: over 60% of the text) of character speech, which generally includes mostly non-events.

### 3.3. Potentials for Literary Studies


Our event types instantiate different degrees of eventfulness, ranging from no (non-event) over little (stative events) to more narrativity (process events and changes of states). Therefore, we could create a simple metric: Non-events are given a narrative value of 0, stative events a value of 2, process events a value of 5 and change of states a value of 7.<sup>6</sup> Based on this metric, we generated narrativity graphs for each manually annotated text.

We did this, as shown in Figure 2, by using cosine weighted smoothing (smoothing window = 50 events). We assume that the narrativity of a section of text results from the narrativity of longer sequences of sentences. This also allows us to evaluate the narrativity graphs, as we will show later.

The graph in Figure 2 shows that we detect highly eventful scenes in the text. Whenever the narrativity value drops sharply in the course of the text, this is due to introspective-reflexive passages.

Additionally, we compare different annotators using narrativity graphs. Figure 3 shows the four parts of Fontane’s *Effi Briest*. Even though there are differences between the two annotators, the graphs are structurally very similar and identify the same passages as peaks

<sup>6</sup>We tested different values without changing the ordering of event types with respect to their narrativity score. Structurally, this seems to have limited impact on the narrativity graphs, as long as said ordering stays intact but further evaluation is required.



**Figure 2:** Narrativity in text course. This example shows the narrativity in Franz Kafka's *Die Verwandlung*. The six most narrative passages are:

1. After the metamorphosis, Gregor exposes himself for the first time to his family and colleague.
2. Gregor leaves his room, his mother loses consciousness, the colleague flees and his father forces him back into his room.
3. Gregor's father throws apples at him. Gregor gets seriously wounded. Escalation of the father-son conflict.
4. Three tenants move into the family's flat.
5. Gregor shows himself to the tenants, who then flee.
6. Gregor dies.

of narrativity.

Against this background, we would consider the automation of event annotation successful if it also allows us to create structurally similar graphs, without necessarily classifying each individual event correctly.

## 4. Automated Event Annotation


Automated event annotation is approached using a two-step process. In the first step, we extract verbal phrases, annotating them with regard to their event type in a second step.

Following the annotation guidelines, we automatically annotate events based on full verbs and their verbal phrases. We rely on a pre-trained tagger and parser [13], finding full verbs in each sentence and, for each verb, walking down the dependency tree to find all tokens they cover. In this tree-walking strategy, we do not descend into relative clauses and stop at conjunctions if their children include full verbs.

Additional rules aim to replicate the annotation guideline's approach to the inclusion of punctuation. For the span evaluation, we exclude all special characters, thereby disregarding errors only related to punctuation or white space. In this scenario, our rule-based tree-walking system yields an F1-Score of 0.71 on *Die Verwandlung* on a per-span basis (meaning each of potentially multiple spans in an annotation is handled individually).

### 4.1. Classifying Event Types from Narrative Theory

Automated classification of event types operates on, sometimes non-continuous, spans of text produced by the first step. Our transformer-based architecture [7] operates on one event candidate generated by the rule-based preprocessor at a time. Individual spans that are part



**Figure 3:** Annotator comparison by narrativity graph on four manually annotated parts in Theodor Fontane's novel *Effi Briest*

of the same event are automatically marked using special tags in the transformer's input vocabulary<sup>7</sup>, thereby allowing the model to focus on the verb of concern if multiple verbs are present in the input sequence. The following example illustrates our input formatting for an event with two text spans.

4. “<EVENT>Vielmehr trieb er</EVENT>, als gäbe es kein Hindernis, <EVENT>Gregor jetzt unter besonderem Lärm vorwärts</EVENT>”<sup>8</sup>

More specifically we rely on an ELECTRA model [6] pre-trained on German data<sup>9</sup>. Our dataset consists of four prose texts, each annotated by a single annotator.<sup>10</sup> In training, we rely on early stopping with a patience of 10 epochs on the F1-scores for all classes weighted by their occurrence, optimising with SGD on a negative log-likelihood loss with label smoothing [26]. Label smoothing was added, together with class weighting for the loss, in an effort to prevent the rare classes from never being predicted.

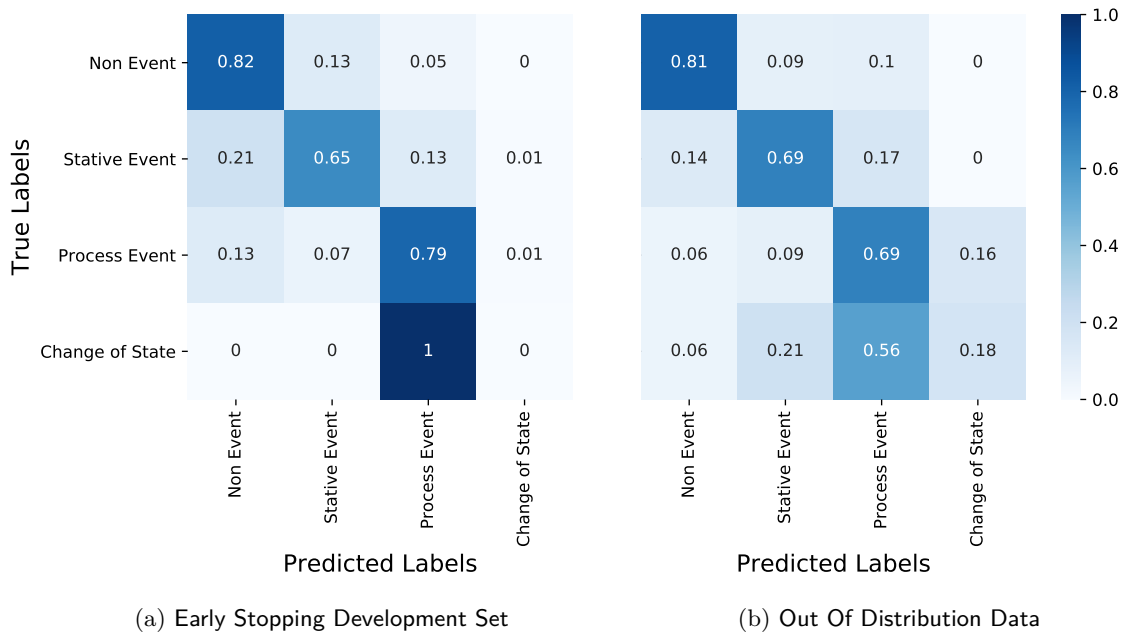
Figure 4a illustrates good performance in recognizing all but one of the classes. Our model performs poorly in distinguishing the process and change-of-state classes. It is important to note that in the development set, the change-of-state class is only found in 4 of 732 events. This can, in part, be attributed to the small number of training examples for the change-of-state

<sup>7</sup>The special tags are placed based on the output of our rule-based processing of the parse tree or, in the case of the classification evaluation, based on spans annotated by human annotators.

<sup>8</sup>“On the contrary, as if there were no obstacle and with a peculiar noise, he now drove Gregor forwards.”

<sup>9</sup><https://huggingface.co/german-nlp-group/electra-base-german-uncased>

<sup>10</sup>For our initial experiments, we use a 80-10-10 train/development/test split while a 90/10 train/development split with a held-out document for testing is employed in the out-of-distribution setup.



**Figure 4:** Confusion matrices using human-annotated spans with the transformer-based classifier for all four classes with row normalized values

class. As a result of the rarity of the change of state class, and given that its members are often attributed to the class with the second-highest narrativity score (limiting the impact on the smoothed narrativity graph), we do not see this result as detrimental to our efforts. Depending on the specific goal, further improvements could be made. Specifically, loss weighting by each class’s narrativity score might improve the model with regards to their application to narrativity graphs. If a more balanced performance for each individual class was desired, early stopping on macro average F1-scores rather than weighted averages would also be an option.

## 4.2. Application to Unseen Documents


While we have established sufficient performance in classifying events from the training distribution, the application to unseen literary works comes with various challenges, including stylistic differences and a change in the distribution of event types.

To validate the model’s performance on such data, we hold out the text *Die Verwandlung* during training, performing early stopping on a subset of events from the remaining documents. Accordingly, the resulting classification performance is broadly applicable to many documents.

While Figure 4b shows no generalization error to out of distribution data when compared to Figure 4a, a limitation of this evaluation remains the impact of propagated errors from the initial span detection step.

To assess the effectiveness of our classification pipeline with regard to narrative graphs, we compare the annotations created by the out-of-distribution model, with those produced by the original annotators. The central characteristics of the graph based on annotations are also to be found in our automatically created graph (see Figure 5). In comparison to the graphs in 5 it may appear as though the model outperforms human annotators. While this can in





**Figure 5:** Comparison of manual and automated annotations by narrativity graph

part be attributed to our system’s good performance, we hypothesize that *Die Verwandlung* is inherently easier to annotate, a theory supported by the reports of our annotators. As discussed previously, we reach an F1-Score of 0.71 in selecting the correct spans; our classification reaches an F1-Score of 0.78 but due to the heavy smoothing employed in the graph, individual errors may not have a large impact on our results. The comparison of our model’s output with manually annotated data gives us confidence in the general applicability of the model to unseen data, justifying its application in corpus analysis.

## 5. Example Corpus Analysis by Automated Event Annotations

We used the trained model to annotate the 2528 texts in the d-Prose corpus [11]. This corpus includes German prose texts from 1870 to 1920.<sup>11</sup> For each text’s respective narrativity graph we extracted the following properties:

- The event count (EC)
- The count of peaks above a relative threshold (PF)<sup>12</sup>
- The peak width at the height of the threshold (PW)<sup>13</sup>
- The average of all peak widths (APW)
- The peak positions (PP)<sup>14</sup>

For illustration, Figure 6 shows the peak properties for the text beginnings of *Krambambouli* and *Das Erdbeben in Chili*.

### 5.1. Text Length and Peak Width


With the first evaluation of the narrativity graphs in Figure 7, we test the influence of text length on the width of the narrative peaks. We are interested in seeing whether longer texts are

<sup>11</sup>To compare the narrativity graphs we first tested dynamic time warping (DTW), a technique originally proposed for speech recognition [1]. Because the implementation of DTW makes length normalizing of our graphs and additional smoothing necessary, we decided to follow another approach.


<sup>12</sup>We used 4 different relative thresholds. Whenever the narrativity graph exceeds 60, 70, 80, or 90 % of the text’s maximum narrativity value, it is considered to be the beginning of a new peak.

<sup>13</sup>The width is the count of events included in a single peak.

<sup>14</sup>The index of the highest value in a peak defines its position.



**Figure 6:** Two examples for peak height, peak width and peak positions with a relative narrativity threshold of 0.8 maximal narrativity value.



**Figure 7:** Text Length (Event Frequency) and Peak Width.

associated with larger peak widths, which from a narratological perspective would be expected. In longer texts, individual passages are more detailed, accordingly one would expect them to contain longer narrative passages.

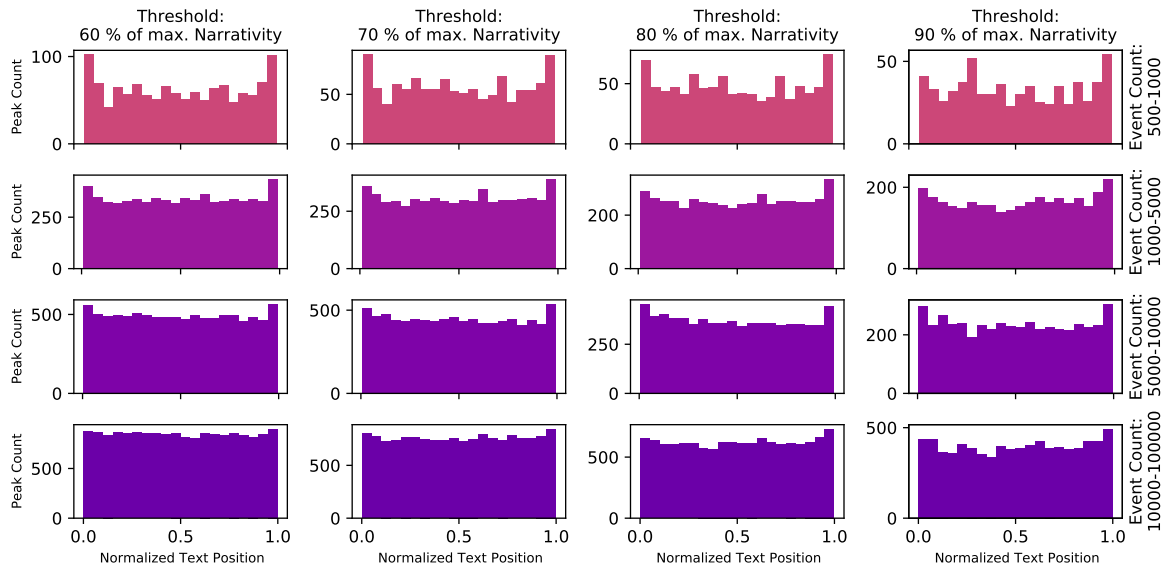
However, the average width of the peaks is not influenced by the length of the texts (correlation for all thresholds  $r < 0.2$ ). This indicates that the alternation of strongly and weakly eventful passages follows similar principles in long and short forms. Only at a threshold of 4 there is a stronger correlation ( $r = 0.19$ ) between the number of events in a text and the average peak width.

## 5.2. Peak Positions and Narrativity

To check which text parts are particularly narrative, we analyze the position of peaks for each of the four thresholds in Figure 8. For that purpose, we divided the corpus into four subcorpora based on each text's event count:

- 339 texts with more than 500 and less than 1,000 events.<sup>15</sup>
- 634 texts with more than 1,000 and less than 5,000 events.
- 329 texts with more than 5,000 and less than 10,000 events.

<sup>15</sup>Due to smoothing in the generation of the narrativity graphs, we cannot examine shorter text with our method.



**Figure 8:** Peak Positions in normalized Text Course. The text course is divided into 20 bins. For each bin the number of peaks over the relative threshold is counted.

- 258 texts with more than 10,000 and less than 100,000 events.

Figure 8 shows a tendency for peaks in all subcorpora to occur at both the beginning and the end.<sup>16</sup> If we compare this with the narrative progression of a long text, such as Fontane’s *Effi Briest*, the high number of peaks in the first part of the text is particularly surprising. We expect longer texts to typically start with descriptive expositions introducing the overall setting and characters, resulting in passages with comparatively little narrativity.

These exploratory analyses show the potential of automatic event classification in conjunction with narrativity graphs for the analysis of large literary corpora. In the future, we plan to integrate the narrativity graphs into text clustering procedures to identify time-, genre- or author-specific text properties.

## Acknowledgments

This work was supported by the DFG through the project “Evaluating Events in Narrative Theory (EvENT)” (grants BI 1544/11-1 and GI 1105/3-1) as part of the priority program “Computational Literary Studies (CLS)” (SPP 2207). We thank Gina Maria Sachse and Michael Weiland for their annotation work.

## References

- [1] D. J. Berndt and J. Clifford. “Using dynamic time warping to find patterns in time series”. In: *KDD workshop*. Vol. 10. 16. Seattle, Washington, USA, 1994, pp. 359–370.

<sup>16</sup>In this context, it is noteworthy that the narrativity of a text’s beginning is highly correlated with its overall average narrativity. (For each the four subcorpora  $r > 0.64$ ).

- [2] J. Björne and T. Salakoski. “Generalizing Biomedical Event Extraction”. In: *Proceedings of BioNLP Shared Task 2011 Workshop*. Portland, Oregon, USA: Association for Computational Linguistics, 2011, pp. 183–191. URL: <https://www.aclweb.org/anthology/W11-1828>.
- [3] X. Carreras and L. Màrquez. “Introduction to the CoNLL-2005 Shared Task: Semantic Role Labeling”. In: *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*. Ann Arbor, Michigan, USA: Association for Computational Linguistics, 2005, pp. 152–164. URL: <https://aclanthology.org/W05-0620>.
- [4] S. Chatman. *Story and discourse: Narrative structure in fiction and film*. Paperback print., [reprint]. Cornell paperbacks. Ithaca, NY: Cornell Univ. Press, 2000.
- [5] N. Chinchor. “MUC-3 Evaluation Metrics”. In: *Proceedings of the 3rd Conference on Message Understanding. Muc3 '91*. San Diego, California, USA: Association for Computational Linguistics, 1991, pp. 17–24. DOI: 10.3115/1071958.1071961. URL: <https://doi.org/10.3115/1071958.1071961>.
- [6] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning. “ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators”. In: *International Conference on Learning Representations*. Online, 2020. URL: <https://openreview.net/forum?id=r1xMH1BtvB>.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota, USA: Association for Computational Linguistics, 2019, pp. 4171–4186. DOI: 10.18653/v1/N19-1423. URL: <https://www.aclweb.org/anthology/N19-1423>.
- [8] G. R. Doddington, A. Mitchell, M. A. Przybocki, L. A. Ramshaw, S. M. Strassel, and R. M. Weischedel. “The automatic content extraction (ACE) program-tasks, data, and evaluation”. In: *Proceedings of LREC*. Vol. 2. 1. Lisbon, Portugal. 2004, pp. 837–840.
- [9] C. J. Fillmore. “Frame semantics”. In: *Linguistics in the Morning Calm*. Seoul, South Korea, 1982, pp. 111–137. URL: <http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199544004.001.0001/oxfordhb-9780199544004-e-013>.
- [10] E. Gius. *Erzählen über Konflikte: Ein Beitrag zur digitalen Narratologie*. Vol. 46. Narratologia, contributions to narrative theory. Berlin and Boston: De Gruyter, 2015. DOI: 10.1515/9783110422405.
- [11] E. Gius, S. Guhr, and B. Adelman. *d-Prose 1870-1920*. Version 2.0. Zenodo, 2021. DOI: 10.5281/zenodo.5015008. URL: <https://doi.org/10.5281/zenodo.5015008>.
- [12] F. Hogenboom, F. Frasinca, U. Kaymak, F. de Jong, and E. Caron. “A survey of event extraction methods from text for decision support systems”. In: *Decision Support Systems* 85.C (2016), pp. 12–22. DOI: 10.1016/j.dss.2016.02.006. URL: <https://research.tilburguniversity.edu/en/publications/a-survey-of-event-extraction-methods-from-text-for-decision-suppo>.
- [13] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd. *spaCy: Industrial-strength Natural Language Processing in Python*. 2020. DOI: 10.5281/zenodo.1212303. URL: <https://doi.org/10.5281/zenodo.1212303>.

- [14] P. Hühn. “Event and Eventfulness”. In: *Handbook of narratology*. Ed. by P. Hühn, J. C. Meister, J. Pier, and W. Schmid. De Gruyter Handbook. Berlin: De Gruyter, 2014, pp. 159–158.
- [15] F. Kafka. *The Metamorphosis and other stories*. New York: Barnes & Noble Books, 1996.
- [16] S. Lahn and J. C. Meister, eds. *Einführung in die Erzähltextanalyse*. 3., aktualisierte und erweiterte Auflage. Stuttgart: J.B. Metzler, 2016. DOI: 10.1007/978-3-476-05415-9.
- [17] Q. Le Sceller, E. B. Karbab, M. Debbabi, and F. Iqbal. “SONAR: Automatic Detection of Cyber Security Events over the Twitter Stream”. In: *Proceedings of the 12th International Conference on Availability, Reliability and Security*. Ares ’17. Reggio Calabria, Italy: Association for Computing Machinery, 2017. DOI: 10.1145/3098954.3098992. URL: <https://doi.org/10.1145/3098954.3098992>.
- [18] M. Martínez and M. Scheffel. *Einführung in die Erzähltheorie*. 10th ed. C.H. Beck Studium. München: C.H.Beck, 2016.
- [19] J. C. Meister. *Computing Action*. Vol. 2. Berlin, New York: Walter de Gruyter, 2003. DOI: 10.1515/9783110201796.
- [20] G. Prince. “A Commentary: Constants and Variables of Narratology”. In: *Narrative* 9.2 (2001), pp. 230–233.
- [21] G. Prince. *A grammar of stories: An introduction*. Vol. 13. De proprietatibus litterarum Series minor. The Hague: Mouton, 2010.
- [22] M.-L. Ryan. “Embedded Narratives and Tellability”. In: *Style* 20.3 (1986), pp. 319–340.
- [23] M.-L. Ryan. *Possible worlds, artificial intelligence, and narrative theory*. Bloomington, Ind.: Indiana Univ. Press, 1991.
- [24] W. Schmid. *Elemente der Narratologie*. 3., erw. und überarb. Aufl. De Gruyter Studium. Berlin: De Gruyter, 2014. DOI: 10.1515/9783110350975.
- [25] M. Sims, J. H. Park, and D. Bamman. “Literary Event Detection”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 3623–3634. DOI: 10.18653/v1/P19-1353.
- [26] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. “Rethinking the Inception Architecture for Computer Vision”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, Nevada, USA, 2016, pp. 2818–2826. DOI: 10.1109/cvpr.2016.308.
- [27] C. Walker, S. Strassel, J. Medero, and K. Maeda. “ACE 2005 Multilingual Training Corpus”. In: *Linguistic Data Consortium, Philadelphia* 57 (2006), p. 45.