‘Orpheus Came to His End by Being Struck by a Thunderbolt’[1]: Annotating Events in Mythological Sequences

Franziska Pannach
Institute of Computer Science
University of Göttingen
franziska.pannach@uni-goettingen.de

Abstract

The mythological domain has various ways of expressing events and background knowledge. Using data extracted according to the hylistic approach (Zgoll, 2019), we annotated a data set of 6315 German sentences from various mythological contexts and geographical origins, like Ancient Greece and Rome or Mesopotamia, into four categories: *single-point* events (e.g. actions), *durative-constant* (background knowledge, continuous states), *durative-initial*, and *durative-resultative*. This data is used to train a classifier, which is able to reliably distinguish event types.

1 Introduction

In narratological terms, *events* have been defined as “constitutive features of narrativity” (Hühn, 2014), the atomic building blocks of a story. An utterance is an event if it communicates a change of state, a “transformation”, which is a fundamental property of any event. In order to produce a plot or a story, events need to follow a chronological or diegetical order, with events being subject to a change in time. *Succession* and *transformation* are therefore key principles in a narrative (Todorov, 1971).

Pustejovsky (2021) distinguished two types of event structures in texts: the surface structure, represented by verbal predicates, and the latent event structure, which refers to sub-events and their representations.

According to Herman (2005), events are often conjoined with *states*, in the sense that a source state S occurs before the transition into a target state S’, triggered by an event (or series of events) E.

In narrative annotation studies, distinctions between events and states are most commonly attributed to the eventfulness of the predicate. The focus on the question what constitutes an event is very much on the question “Who does what to whom?”

This works presents the annotation efforts to classify different types of events in the mythological and religious domain. Textual sources from those domains often do not narrate plots in a straightforward manner. Instead, they use stylistic devices, like prolepses, to transport their narrative, which can make automatic extraction and event labelling challenging.

For this study, different types of events and their chronological order have been extracted by domain experts from the fields of Ancient Near Eastern Studies, Religious Studies and Classics from a large variety of sources.

For each source, a sequence of events and background information was manually extracted based on the original, e.g. in ancient Greek or Sumerian, where available. Those sequences were derived according to the hylistic approach (Zgoll, 2019) from the narrative domains of mythological and religous studies (Zgoll and Zgoll, 2020; Gabriel et al., 2021). Each sequence corresponds to one variant of a (mythological) plot in the respective source.

The context window of the myth variant, i.e. which passages of the source refer to a mythological plot, is identified by the domain expert. Hence, the text passages that correspond to the sequences, as well as the sequences themselves, differ in length. The sequences can be used for comparatistic tasks, such as the comparison of narrative plots or background information, e.g. the characterization of entities. The distribution of disciplines from which the sources are taken are presented in Figure 2.

The narrative sequences contain practically no discourse markers, and are comprised of individual statements (*hylemes*) which are strictly in present tense. The hylemes contain events, including stative events or states in the chronological (narrative) order, not the diegetic order.

As an example, we use the quote ‘Orpheus

Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII), pages 10–18

came to his end by being struck by a thunderbolt'[1]. From this sentence, the following short sequence of statements (*hyleme sequence*) can be manually extracted:

1. 'Orpheus is struck by a thunderbolt.'
2. 'Orpheus dies.'
3. 'Orpheus is dead.'

The task of this work is to annotate the *event* types of each individual statement (*hyleme*). The data is in German, examples have been translated by the author for this paper, where necessary.

This paper is structured as follows: Section 2 frames the work into context of similar annotation efforts. Section 3 compares the categorisation of events used in this work with previous annotation efforts by Gius and Vauth (2022). The data set used for this study is described in Section 4. In Section 5, we describe the annotation effort and its results. We present a simple classifier to determine the event types, which we describe in Section 6. Finally, the paper ends in a discussion in Section 7.

2 Related Works

Our work is situated in the context of mythological research, but has potential for application in other domains. On the linguistic level, it is related to the study of lexical aspect (or *Aktionsart*) and the situation entity (SE) annotation task. Friedrich et al. (2016) label SE types from clauses in a supervised sequence modelling task using features of the main verb, its main referent, and the clause itself. They report good results across different genres. In an earlier study, Friedrich and Pinkal (2015) annotated clausal aspect for automatically recognising whether a clause describes a *habitual*, *episodic*, or *static* phenomenon.

Metheniti et al. (2022) sucessfully identified temporal aspect (*telicity* and *duration*) in English and French data sets using a transformer approach.

Furthermore, there are a number of practical approaches which attempt to define and narrow the narrative concept of *events* and their representation.

Chambers and Jurafsky (2008) introduce an approach to use unsupervised learning of event chains centered around an event protagonist. They train a temporal classifier to produce a temporally ordered narrative chain. In a subsequent study, they present the concept of narrative schemas, as "coherent sequences or sets of events" (Chambers and Jurafsky, 2009). By applying an unsupervised learning approach, they add semantic roles to the argument structure of their event chains. Multiple events chains are then combined into a narrative schema.

TimeML (Pustejovsky et al., 2005a,b) is a markup language designed based on XML which provides a standardized way of annotating temporal expressions and events in text, including the temporal relationships between events. It is used for the annotation of temporal and event information. Four automatic TimeML annotation systems have been evaluated by Ocal et al. (2022).

Reiter (2015) compared the annotation of narrative segments performed through crowd-sourcing, by student annotators and summary annotations. Kwong (2011) annotated a corpus of fables regarding their structural and semantic properties, including temporal information. Events that are part of a script, such as 'baking a cake', have been automatically mapped to narrative texts by Ostermann et al. (2017).

Events and event types in narrative plots have been studied by Gius and Vauth (2022). They operationalize the concepts of narrativity and tellability as discourse phenomena. They use spans of text defined by finite verbs as annotation units. Guis' and Vauth's concepts of states and events are probably closest to those of the hylistic theory presented by (Zgoll, 2019). Therefore, we will compare the two annotation approaches in more detail in the next section.

3 Event Categories

The narrative event model of Gius and Vauth (2022) and Vauth and Gius (2021) uses four categories of events: *change of state*, *process events*, *stative events*, and *non-events*. The basis of their event representation is the finite verb in 'minimal sentences', i.e. all tokens that are assigned to the verb.

In contrast, the categories used for hylistic analysis are: *single-point* (punctual), *durative-constant*, *durative-initial* and *durative-resultative*. We classify a statement (*hyleme*) into one of these four categories, but the value is of course mainly associated with the verb. In both theories, each annotation unit has one finite verb. Figure 1 illustrates the difference between four hyleme types.

Single-point hylemes are true at one point during the narrative sequence extracted from the source. This includes active actions, passive experiences,

[1] Pausanias, Description of Greece

reactions, perceptions or feelings. The *single-point* event has its beginning and end during the sequence. However, that does not necessarily indicate an event with a short duration.

Durative hylemes hold true for a part of the sequence or over the course of the entire sequence. There are three sub-types: *Durative-constant*, *durative-initial*, and *durative-resultative*. *Durative-constant* hylemes are always true, e.g. "Orpheus is the son of Oeagrus." They often communicate background knowledge about the narrative. Additionally, certain 1N/nS statements[2] (Genette, 1983), e.g. "Hades works the sails" are also considered *durative-constant*.

There are two types of states which are true over a part of the sequence, but change their value at some point. *Durative-initial* hylemes are true at the beginning of the sequence. *Durative-resultative* hylemes are statements that become true at some point during the sequence (e.g. 'Orpheus is dead.') and remain true for the rest of the sequence. In the mythological domain, these context-sensitive hylemes often connect contexts and plots. Hyleme sequences follow a relative temporal order, without discourse markers.

Table 1 shows how different example sentences from Kafka's *Metamorphosis* are annotated according to both theories.

Guis and Vauth's category 'change of state', used for the first example sentence in Table 1, corresponds widely to the *single-point* category that is used for the annotations presented in this work. However, the category 'change of state' can be realised with different properties (Gius and Vauth, 2022). One of those properties is *iterative*. In most cases where this property would be applied, the hylistic theory would dictate the annotation of *durative-constant* (resp. *-initial* or *-resultative*), e.g. "Charon works the sails". This statement refers to an action that is characteristic for a character. It can be either ongoing, continuous, or characteristic in the sense that *Charon* is someone who is capable of performing this action.

The second example sentence "found he himself in his bed into a monstrous insect-like creature transformed" would be annotated as *single-point* statement according to hylistic theory, because the predicate "found" implies that he realises he has been transformed into a bug exactly once[3] during the course of the narrative. However, a hylistic analysis of the plot would necessarily include a statement like "Gregor Samsa is a human transformed into a bug", which would be annotated as *durative-constant*. This statement does not need to be explicitly stated in the text, it can be implied. The sentence "His room lay quietly between the four well-known walls" demonstrates where the main difference between the two theories lie:

According to Gius and Vauth (2022), this sentence is annotated as a *stative-event*. While the hylistic theory (Zgoll, 2019) also recognises that this is an ongoing state, it distinguishes between types of ongoing states. Hylemes that are valid at the beginning, but change during the course of the narrative are categorised as *durative-initial*, e.g. *Eurydice is alive*. Hylemes that are the result of an event, e.g. *A snake bites Eurydice* $\rightarrow$ *Eurydice is dead*, are *durative-resultative*. Thirdly, there are hylemes that are true over the entire course of the narrative, e.g. *Eurydice is Orpheus' wife.* Those statements are *durative-constant*. They communicate the background knowledge that is the basis of a narrative, e.g. information about characters, their relations between each other and properties of the world in which a (mythological) story takes place. In order to determine the hylistic event category of the third sentence, therefore, we need to establish if the quietness of the room is a) the result of something that happened previously, or b) the initial state that is changed later-on, e.g. by someone barging in, or c) a general characterisation of the room. *Durative-resultative* statements are often preceded by a *single-point* statement, which corresponds to a *change of state* event according to Gius and Vauth (2022). However, occasionally *durative-resultative* statements are the result of the entire narrative, e.g. "No one can solve this incantation" is the result of the entire narrative of the invocation MS 2353 (CUSAS 32, 19a) (George, 2016; Rudik, 2011).

Non-events are not represented in hylistic theory, because they do not contain plot relevant information. Non-events contain mainly conditional, subjunctive, or modalised statements (Vauth and Gius, 2021).

4 Data

As explained in the previous section, the event definition used in this paper is different from the ones mentioned in Section 2. Furthermore, event

[2] "narrating one time what happened n times"

[3] Afterwards he knows that he is a bug. (*durative-constant*)

Table 1: Comparison of Event and State Categories

Sentence	Guis and Vauth, 2022	Hylistic Class
"Gregor Samsa one morning from uneasy dreams awoke"	Change of state	single-point
"found he himself in his bed into a monstrous insect-like creature transformed"	Process Events	single-point
"His room lay quietly between the four well-known walls"	Stative Events	durative(-constant, -initial, or -resultative dependent on context)
"She would have closed the door to the apartment"	Non-events	*NA*

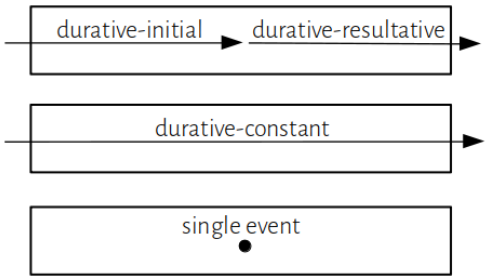

Figure 1: Types of Hylemes

statements (hylemes) are not derived directly from the textual representation in a source. Mythological plots and descriptions of background knowledge are often not told in a straightforward manner. Rather, they allude to related aspects of similar myths, and use comparisons, context and intertextuality in ways that makes the interpretation of what exactly happens in a myth variant hard to understand for laymen and even harder to process using NLP tools. Even the order of events is sometimes difficult to establish, as the following example illustrates:

> (1) "But Orpheus, son of Oeagrus, [they sent back$_4$ with failure from Hades], [showing$_3$ him only a wraith] [of the woman for whom he came$_2$]; [her real self they would not bestow$_3$], [for he was accounted to have gone upon a coward's quest$_1$], ..."[4]

We can see that the sequential order of events is different from the order presented in the source. Chronologically, *Orpheus* first goes on a coward's quests (1), in order to rescue his wife (2), but they (= the inhabitants of the netherworld) do not give him his real wife (= Euydice), but show him only a wraith of her (3). As a result, they send him back with failure (4). Fictional texts often follow their own order or use non-linear narrative, in order to create tension or highlight certain aspects of the plot. In-text annotations can rarely account for the discrepancy, especially if the events are presented without discourse markers or temporal expressions. The hylistic theory distinguishes between the order in the source and the chronological order.

The next example will illustrate how main plot events in classical sources are communicated merely by allusion.

> (2) "If Orpheus, arm'd with his enchanting lyre,
> The ruthless king with pity could inspire,
> And from the shades below redeem his wife;"[5]

In this variant of the myth *Orpheus and Eurydice*, we know that *Orpheus* has a lyre, which has some enchanting properties. He successfully inspires some unnamed ruthless king (possibly *Dis* through *Proserpina* (Bowra, 1952)). Exactly how he achieves this is left out, because this passage might allude to other variants of the myth, where this is discussed in more detail. Then *Orpheus* redeems his wife from the shadows below, alluding to the netherworld (*Hades*). This information alone does not tell us much about what exactly takes place. In Georgics, 4, 453–527, Vergil himself tells a more detailed story of how the events took place. This includes how exactly *Eurydice* dies, and the

[4]Plato Symp. 179d http://data.perseus.org/citations/urn:cts:greekLit:tlg0059.tlg011.perseus-eng1:179d

[5]Vergil. Aeneid. 6, L.98-123 http://data.perseus.org/citations/urn:cts:latinLit:phi0690.phi003.perseus-eng1:6.98-6.123

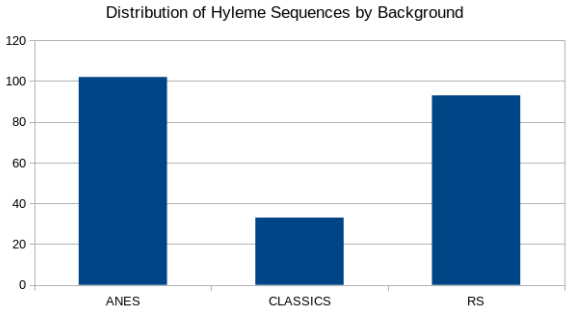

Figure 2: Distribution of narrative sequences (*hyleme sequences*) by topic, Ancient Near Eastern Studies (ANES): 102, Classics: 33, Religious Studies: (RS) 93

fact that *Orpheus* is presented with conditions for bringing his wife back from the netherworld (i.e. he is not allowed to look at her).

Both examples show that extracting information from the texts is a challenging task that needs to be guided by informed scholars. This issue severely magnifies if we do not consider the well documented Classical domain, but extend studies to fields like Ancient Near Eastern Studies, where sources are often scarce, and their supporting material (e.g. cuneiform on stone tablets) can be damaged or difficult to read.

Therefore, the context-window, the plot inherent events and background knowledge presented in the 228 sources have been extracted manually. Each source is presented in one sequence of event statements, so called *hyleme sequence* (Zgoll, 2019). The hylemes were originally not annotated with their state or event types. However, in order to process the sequences for further study using NLP methods, e.g. measuring the similarity of plots or aligning variants of the same myth, the annotation of *single-point* events, and *durative* statements was needed.

Each hyleme sequence describes the plot of one myth variant and related background information. The statements usually do not contain fixed or relative temporal expressions, or relations such as *before* or *after*. Instead, the succession of events is expressed through the sequential order.

The annotated data is a set of 6315 hylemes and their assigned category. It is not, as discussed above, an annotation of concurrent text from the sources, but sequences describing the plot that were extracted manually.

The statements themselves are usually short, concise sentences in German, consisting of only main clauses, containing one finite verb in present tense and active voice (where possible). Co-references are widely avoided. Instead, each statements contains the resolved arguments, which are repeated in the subsequent statements, even if they are only communicated by co-references in the text. One sentence in a source can translate to multiple statements, e.g. “Orpheus is the son of Oeagrus.”, “The gods send Orpheus back from Hades as a failure.”,... Aspects which are alluded but can be safely determined by the informed scholar (e.g. Orpheus’ wife’s name is Eurydice) are added in square brackets. Those implications can be part of the statement, e.g. a name, or an entire statement. For instance, in the example sentences from Kafka’s *Metamorphosis*, the first statement “Gregor Samsa one morning from uneasy dreams awoke” would be preceded by a statement like “Gregor Samsa is sleeping” in a hylistic analysis.

5 Annotation

The data set was annotated by six annotators. Since *durative-initial* and *durative-resultative* statements are context-sensitive, annotators always processed the entire sequence. Each narrative sequence was annotated twice. Table 2 gives an overview of the annotators’ disciplines, and level of education.

Annotator	Background	Level of Education
A1	ANES	B.A.
A2	CS/CL	M.Sc.
A3	Classical Studies	B.A.
A4	ANES/DH	Doctoral Deg.
A5	ANES	B.A.
A6	ANES	M.A.

Table 2: Annotators’ backgrounds

All annotators had previous experience with the hylistic theory. Additionally, they were trained in an initial annotation meeting. Each annotator was given a set of sequences, which were annotated individually and discussed by the group afterwards. Annotators were also given a set of guiding questions and example statements to help them chose the right event category where in doubt. The guiding questions were presented in a flowchart. Additionally, annotators with explicit knowledge in the field, e.g. Classics, were also asked to check the original sources for guidance where in doubt. For example, the English statement “Orpheus brings

back the dead (from the netherworld)" can be interpreted as *single-point* or *durative-constant*. Through the original Greek source, it can be determined that it should be annotated as *durative-constant*, because the imperfect form (ἀνῆγεν) is used (Bowra, 1952). In a second meeting, questions that arose during the annotation process were discussed.

Items were annotated in 11 different pairings, with varying first and second annotators. In all but one cases, the inter-annotator agreement for the annotation task ranges from substantial (κ 0.61-0.80) to almost perfect agreement (κ 0.81-0.99). The agreement is reported in Table 3. Annotator pairs A2-A4 and A4-A5 have perfect agreement over the shared annotations. Pair A2-A5 has a relatively low value of $\kappa = 0.4$. This is due to one particularly long sequence containing 114 hylemes. Many statements in this sequence contain descriptions of a mythical house, e.g. "The vault of the house is a rainbow". These were annotated as *durative-constant* by one annotator, while the other interpreted these descriptions as results of some action in the sequence, and therefore annotated them as *durative-resultative*. Consequently, event type annotations of all descriptions of the house in that sequence are mismatching (consequential error). This results in a low overall κ for the annotator pair A2-A5.

Pair	No. of items	Cohen's κ
A1-A2	4552	0.848930
A1-A3	398	0.874665
A1-A4	299	0.929306
A1-A5	149	0.733025
A1-A6	96	0.631285
A2-A3	187	0.918325
A2-A4	90	1
A2-A5	127	0.402008
A3-A4	136	0.866710
A3-A5	239	0.811959
A4-A5	42	1

Table 3: Inter-annotator agreement (Cohen's κ) between pairs of annotators

In cases where the first and second annotator disagreed, the gold standard was derived by discussion in a separate meeting, or following the judgement of the annotator whose discipline the sequence belongs to. Performance of annotators against gold standard, and total number of annotated items are reported in Table 4.

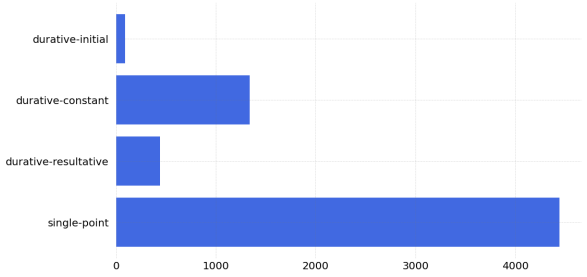

Figure 3: Distribution of the event types in the final data set (gold standard annotation)

Annotator	Gold	No. of items
A1	0.939978	5494
A2	0.914271	4956
A3	0.951389	960
A4	0.953625	567
A5	0.705362	557
A6	0.631285	96

Table 4: Cohen's κ of annotators against Gold standard

The final gold labels are an important foundation for the next analyses, e.g. plot comparison and alignment or the comparison of background information in the individual sources. The distribution of the gold-standard labels is shown in Figure 3. The majority of the data consists of *single-point* statements, of the *durative* statements, the *durative-constant* hylemes are the largest group. Three hylemes had to be excluded from the data, because their types could not be determined (e.g. the statement "The *kur-ĝara* and *gala-tur* ...?" has a missing predicate due to the source not being properly readable).

6 Classifier

Based on the gold labels of the annotation as described in the previous section, two event type classifiers were trained.[6] The resulting models can be used to pre-classify new statements, and to classify statements in future data sets that can be used for comparison, e.g. including movie adaptations of mythological narratives.

The separation of the data into *durative* and *single-point* statements is an important first step for the subsequent analyses of the narrative sequences,

[6]The classifiers and an excerpt of the annotated data can be found under: https://gitlab.gwdg.de/franziska.pannach/hylva_event_types For access to the full data, kindly contact the author.

since *single-point* statements correspond to events, whereas *durative* hylemes correspond to descriptions of background knowledge.

The task to automatically classify event types is not trivial. Especially, automatically distinguishing the three types of *durative* hylemes is challenging. This is due to multiple reasons. Firstly, the three classes are unbalanced, with more *durative-constant* hylemes, and very few *durative-initial* hylemes. Additionally, *durative-initial* or *-resultative* hylemes can be quite similar to *durative-constant* hylemes in terms of vocabulary and grammatical structure. As discussed above, their value is often context-sensitive.

For the classification task, a multinomial naive bayes model was selected. For that purpose, the data set was split into a training and test set with a split of 75 %-25 %. The hyper-parameters were selected by performing a grid search. In particular, the grid search established whether the feature vector is best constructed using a bag-of-words or TF-IDF vectorizer.

As a result, the hyper-parameters were set as: Laplace smoothing parameter $\alpha = 0.01$, bag-of-words features, and an n-gram range of 3.

Firstly, we analyze the results for binary classes *single-point* and *durative*, which combines *durative-initial*, *durative-constant*, and *durative-resultative* statements. For that purpose, all three labels were subsumed under the coarse class *durative* for training. The binary classifier performs well on *single-point* hylemes, and reasonably on *durative* hylemes. The performance of the classifier is reported in Table 5.

Secondly, we investigate how the classifier performs if trained on just the different types *durative* hylemes. For that purpose, all *single-point* hylemes were removed from the training and test set. The majority of the test set consists of *durative-constant* hylemes (69 %) and *durative-resultative* hylemes (24 %). The results are reported in Table 6.

Lastly, we present the classifier for the classification of fine-grained classes. It was trained on the entire training and test set including fine-grained *durative* classes. A second classifier combining the first two models (binary and durative-only) in two steps was trained but did not improve results.

Table 7 shows the performance of the fine-grained classifier. The confusion matrix for the classifier is shown in Figure 4. We can see that the classifier favours the *single-point* class. This is most apparent in the case of *durative-constant* statements, which were misclassified as *single-point* in 70 cases.

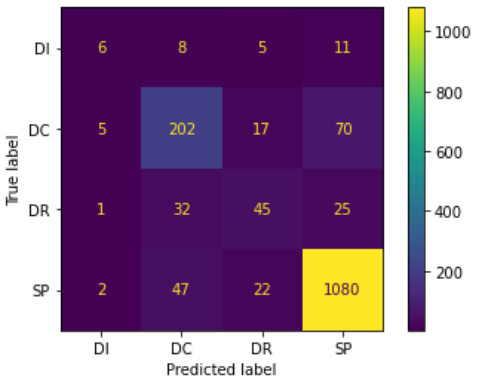

Figure 4: Confusion matrix for the classifier trained on the gold labels, DI = *durative-initial*, *durative-constant*, DR = *durative-resultative*, SP = *single-point*

	Precision	Recall	F1
durative	0.83	0.75	0.79
single-point	0.91	0.94	0.92

Table 5: Performance of the binary classifier

	Precision	Recall	F1
dur.-initial	0.50	0.23	0.32
dur.-constant	0.81	0.90	0.85
dur.-resultative	0.62	0.51	0.56

Table 6: Performance of the durative classifier

7 Discussion

In order to annotate event types for the mythological and religious domains from the source, the sequence of events and background information has to be extracted. Automatically extracting these events from can be challenging, as demonstrated in the examples in Section 4. Therefore, the sequences of statements describing events and states from the sources was achieved manually. Subsequently, we present annotations based on the hylistic theory (Zgoll, 2019), which was developed specifically for the mythological domain, but can be easily applied to other types of narrative as well. The data includes over 6300 statements from 228 narrative sequences. The statements have been annotated into four categories. *Single-point* statements, communicating events, *durative-constant* (background information), *durative-initial*

	Prec.	Recall	F1	Support
d.-initial	0.50	0.17	0.25	30
d.-constant	0.72	0.67	0.69	294
d.-resultative	0.55	0.45	0.49	103
single-point	0.90	0.95	0.93	1151

Table 7: Performance of the fine-grained classifier

and *durative-resultative*, which hylemes indicate that their truth value changes during the course of the sequence. After training the annotators, an overall satisfying inter-annotator agreement κ was reached.

The main weakness of the presented approach is that the event categories are not assigned directly to the text. This is due to the original source material being extremely diverse in form, language, and genre. Instead, the labels are assigned to the *hyleme sequences* which require significant manual effort and knowledge of the original material.

Durative labels, especially *durative-initial* and *durative-resultative*, are context-sensitive. The value of a statement has to be assessed within the context of the narrative sequence. Since two identical statements can have different labels in different contexts, the classification task is particularly challenging. This is the case especially if the label depends not only on a single preceding statement (e.g. *Eurydice dies.* $\rightarrow$ *Eurydice is dead.*), but on the entire sequence (e.g. *Nobody can solve this invocation.*)

When hyleme sequences are extracted from modern texts in well-resourced languages, such as German or English, the manual effort could be alleviated by employing NLP methods, such as named entity recognition or semantic role labelling. With a larger number of texts and corresponding sequences, it would also be possible to automatically identify candidate statements from text.

The gold standard data represents the actual distribution of labels, i.e. *single-point* statements (actions) are more prevalent than *durative* statements. Hence, the final data set is skewed which explains the performance of the classifier. In this work, a simple Naive Bayes classifier was implemented for demonstration purposes. A more sophisticated model, e.g. following a multi-lingual transformer approach (Conneau et al., 2020), would potentially deliver better results.

In future studies, the plots of mythological and religious narrative can now be studied and compared using NLP and alignment techniques on sequences of *single-point* statements. The background information in *durative-constant* can be included, or processed separately to represent the narrative-inherent background knowledge.

References

C. M. Bowra. 1952. Orpheus and Eurydice. *The Classical Quarterly*, 2(3-4):113–126.

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.

Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Annemarie Friedrich, Alexis Palmer, and Manfred Pinkal. 2016. Situation entity types: automatic classification of clause-level aspect. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1757–1768, Berlin, Germany. Association for Computational Linguistics.

Annemarie Friedrich and Manfred Pinkal. 2015. Automatic recognition of habituals: a three-way classification of clausal aspect. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2471–2481, Lisbon, Portugal. Association for Computational Linguistics.

Gösta Ingvar Gabriel, Brit Kärger, Annette Zgoll, and Christian Zgoll, editors. 2021. *Was vom Himmel kommt:Stoffanalytische Zugänge zu antiken Mythen aus Mesopotamien, Ägypten, Griechenland und Rom*. De Gruyter, Berlin, Boston.

Gérard Genette. 1983. *Narrative discourse: An essay in method*, volume 3. Cornell University Press.

Andrew R George. 2016. *Mesopotamian incantations and related texts in the Schøyen Collection*. CUSAS: Cornell University Studies in Assyriology and Sumerology 32.

Evelyn Gius and Michael Vauth. 2022. Towards an event based plot model. a computational narratology approach. *Journal of Computational Literary Studies*, 1(1).

David Herman. 2005. Events and event-types. *Routledge Encyclopedia of Narrative Theory. London: Routledge*, pages 151–52.

Peter Hühn. 2014. Event and eventfulness. In Peter Hühn, John Pier, Wolf Schmid, and Jörg Schönert, editors, *the living handbook of narratology*. Hamburg University, Hamburg.

Olivia OY Kwong. 2011. Annotating the structure and semantics of fables. In *Proceedings of the 25th Pacific Asia Conference on Language, Information and Computation*, pages 275–282.

Eleni Metheniti, Tim Van De Cruys, and Nabil Hathout. 2022. About time: Do transformers learn temporal verbal aspect? In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 88–101, Dublin, Ireland. Association for Computational Linguistics.

Mustafa Ocal, Adrian Perez, Antonela Radas, and Mark Finlayson. 2022. Holistic evaluation of automatic timeML annotators. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1444–1453.

Simon Ostermann, Michael Roth, Stefan Thater, and Manfred Pinkal. 2017. Aligning script events with narrative texts. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 128–134, Vancouver, Canada. Association for Computational Linguistics.

James Pustejovsky. 2021. The role of event-based representations and reasoning in language. In Tommaso Caselli, Eduard Hovy, Martha Palmer, and Piek Vossen, editors, *Computational Analysis of Storylines: Making Sense of Events*, chapter 1, pages 23–46. Cambridge University Press.

James Pustejovsky, Robert Ingria, Roser Saurí, José M. Castaño, Jessica Littman, Robert J. Gaizauskas, Andrea Setzer, Graham Katz, and Inderjeet Mani. 2005a. The specification language timeml. In *The Language of Time - A Reader*.

James Pustejovsky, Robert Knippen, Jessica Littman, and Roser Saurí. 2005b. Temporal and event information in natural language text. *Language resources and evaluation*, 39:123–164.

Nils Reiter. 2015. Towards annotating narrative segments. In *Proceedings of the 9th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH)*, pages 34–38, Beijing, China. Association for Computational Linguistics.

Nadezda Rudik. 2011. *Die Entwicklung der keilschriftlichen sumerischen Beschwörungsliteratur von den Anfängen bis zur Ur III-Zeit*. Ph.D. thesis, Friedrich-Schiller-Universität Jena.

Tzvetan Todorov. 1971. The 2 principles of narrative. *Diacritics*, pages 37–44.

Michael Vauth and Evelyn Gius. 2021. Richtlinien für die Annotation narratologischer Ereigniskonzepte.

Annette Zgoll and Christian Zgoll, editors. 2020. *Mythische Sphärenwechsel: Methodisch neue Zugänge zu antiken Mythen in Orient und Okzident*. De Gruyter, Berlin, Boston.

Christian Zgoll. 2019. *Tractatus mythologicus: Theorie und Methodik zur Erforschung von Mythen als Grundlegung einer allgemeinen, transmedialen und komparatistischen Stoffwissenschaft*. De Gruyter, Berlin, Boston.