Fables-DTR: A Corpus of Fables Annotated for Discourse and Temporal Relations

Purificação Silvano[1], António Leal[1], Maciej Ogrodniczuk[2], Martyna Lewandowska[2], Anna Śliwicka[2], Aleksandra Tomaszewska[2], Luís Filipe Cunha[3], Evelin Amorim[3], Joana Gomes[1], Alípio Jorge[4]
[1] Faculty of Arts and Humanities (FLUP), University of Porto, Portugal
[2] Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland
[3] INESC TEC, Porto, Portugal
[4] Faculty of Sciences (FCUP), University of Porto, Portugal
{msilvano,jleal}@letras.up.pt, m.ogrodniczuk@ipipan.waw.pl, martyna.lewandowska.ut@gmail.com
anna.sliwicka@interia.pl, aleksandratomaszewska15@gmail.com, {luis.f.cunha,evelin.f.amorim}@inesctec.pt
joana.gomes2003@hotmail.com, amjorge@fc.up.pt

Abstract

This paper presents Fables-DTR, a corpus of Aesop's fables annotated for discourse and temporal relations, designed to explore how event sequencing and aspectual features and discourse relations interact. Building on the ISO 24617 Semantic Annotation Framework, integrating Part 1 (Time and Events) and Part 8 (Discourse Relations), the resource provides a unified representation of discourse structure and temporal and aspectual features. The corpus comprises 15 fables in English, automatically translated into European Portuguese and Polish (45 texts in total), with all translations manually validated by native linguists to preserve semantic and discourse features. Each fable is annotated in two layers: (i) for discourse relations, argument roles, and signals; (ii) for temporal relations, and event attributes, such as Tense, Aspect, Polarity. The resulting dataset provides relevant information about the association between discourse relations and their temporal and aspectual features. Fables-DTR contributes both a valuable resource for cross-linguistic and narrative discourse analysis and empirical evidence for integrating ISO standards in multilayer annotation. It also provides a foundation for computational applications in discourse parsing, event ordering, and implicit relation detection.

Keywords: discourse relations, temporal and aspectual features, implicit and explicit discourse relations, annotated dataset, ISO-based annotations

1. Introduction

Coherence in discourse is shaped by many factors, but it is significantly influenced by the interaction between semantic relations and temporal structure. Readers interpret events not only through explicit connectives but also through implicit cues of order, causality, and aspect. Discourse frameworks have long aimed to capture these, yet empirical work linking discourse and temporal annotation remains scarce.

Research on discourse relations and temporal relations relies on annotated data, but the available corpora are narrow in both genre and language diversity. Moreover, the existing datasets rarely follow a unified annotation standard, which makes cross-linguistic or cross-theoretical comparison difficult. Although the ISO 24617 Semantic Annotation Framework (semAF) offers interoperable specifications for temporal and discourse relations, its potential for joint, multilingual analysis remains largely unexplored.

This paper addresses these limitations by presenting Fables-DTR, a resource that connects discourse structure and temporal and aspectual features. The corpus, composed of Aesop's fables, is annotated with discourse relations and temporal and aspectual information following an annotation scheme that is integrated and interoperable. Using Aesop's fables as a controlled setting allows us to study how event sequencing and coherence operate in short, self-contained narratives that are rich in discourse relations.

Beyond corpus creation, this work aims to contribute methodological evidence on how discourse-temporal annotation can be harmonized. It also provides empirical grounding for tasks that require understanding beyond the sentence level, such as implicit relation detection, narrative event ordering, and multilingual discourse parsing. Fables-DTR strengthens the empirical foundation for research on how time and discourse jointly structure meaning. The following sections situate the study in prior work (Section 2), describe corpus construction and annotation principles (Sections 3 and 4), and present analyses that highlight tendencies about the relation between the discourse and temporal structure (Section 5).

2. Related work

Textual coherence depends on semantic links among discourse segments. These links, discourse relations (DRs), are formalized in frame-

Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026), pages 1856–1868
11–16 May 2026.

works such as Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) and Segmented Discourse Representation Theory (SDRT) (Asher and Lascarides, 2003). RST models texts as hierarchies of nucleus–satellite relations grounded in communicative function, whereas SDRT uses formal semantics to integrate discourse coherence with inference. Both frameworks account for phenomena such as temporal anaphora and global discourse organization, underscoring the role of implicit relations that lack overt connectives. Building on this tradition, prior research on adverbial perfect participial clauses has shown that factors such as linear order, temporal sequencing, and aspectual class predict specific DRs and that these parameters vary cross-linguistically, motivating the integrated treatment of discourse and temporal and aspectual features pursued in this work (Silvano et al., 2023).

Resources that jointly annotate discourse and temporal layers remain scarce. Kaneko (2016) introduces a Japanese corpus that captures temporal, discourse, and multilayered situational relations across 170 Wikinews fragments. The Groningen Meaning Bank (Bos et al., 2017) constitutes a multilayer semantic annotation of public-domain texts, representing discourse relations and tense information within the Discourse Representation Theory (DRT) framework (Kamp and Reyle, 1993). In this resource, temporal relations are not explicitly annotated but are instead inferred from the formal representation of situations and their temporal and aspectual properties. Datasets combining temporal and causal relations are comparatively more common. Two prominent examples include the Causal-TimeBank (Mirza et al., 2014), which extends the TimeML framework by automatically annotating explicit causal relations, and the Event StoryLine Corpus (Caselli and Vossen, 2017), a multilayer annotation of news stories that focuses on when events occurred and why they happened, also employing TimeML-based notations.

In addition to the scarcity of annotated datasets that jointly capture discourse and temporal relations, proposals for unified annotation frameworks encompassing both layers remain rare. The ISO 24617 – Language Resource Management: Semantic Annotation Framework (SemAF) defines a modular semantic architecture with several interoperable components, namely referential, temporal, and spatial, designed to promote language-agnostic and cross-domain interoperability. Despite its comprehensive design, the framework has seen limited practical adoption, with the exception of Part 1 (ISO, 2012), which formalizes temporal and event annotation (e.g., Pustejovsky et al., 2006). With respect to discourse relations, ISO 24617-8 (ISO, 2016b) provides a language- and genre-independent scheme for annotating local discourse relations, including systematic mappings to major theoretical frameworks, yet it too remains underexplored. Recent multilingual implementations, covering English, Polish, and European Portuguese, demonstrate its suitability for comparative discourse research (e.g., Silvano et al., 2023, Tomaszewska et al., 2024, Ogrodniczuk et al., forthcoming). However, despite the crucial role of temporal features in shaping discourse inferences, to the best of our knowledge, no annotation scheme has yet integrated both Parts 1 and 8 of the ISO 24617 framework into a unified model.

This paper addresses these limitations by introducing Fables-DTR, a corpus of fables annotated for discourse and temporal relations. This resource integrates discourse structure with temporal and aspectual features within a unified framework, offering new insights into how narrative structure, temporal features, and aspect jointly contribute to textual coherence.

3. Fables-DTR

To conduct a comparative analysis of discourse relations across typologically distinct languages, we compiled a specialized corpus of Aesop's fables in English, European Portuguese, and Polish. In the next subsection, we present the details of this dataset and the process of its development.

3.1. Dataset creation

The initial source texts were 45 Aesop's fables in English retrieved from a single source, the the Standard Ebooks website[1], based on the popularity of Aesop's fables according to Wikipedia.

To ensure structural parallelism for cross-linguistic comparison, the English fables were automatically translated into the other two languages using DeepL[2]. After careful consideration and several attempts to pursue alternative approaches, we concluded that translating the fables from a single source text was the most appropriate solution. Existing translations from the original Ancient Greek were available in both Portuguese and Polish; however, comparison of the versions in the three languages revealed substantial differences in structure, lexicon, and syntax. Moreover, producing new translations directly from Ancient Greek would have required a level of expertise that the authors do not possess. For these reasons, English was selected as the source language. Since our objective was to construct a parallel corpus that would enable systematic comparison of how discourse and temporal structure are realized across languages, we opted

[1] https://standardebooks.org/ebooks/aesop/fables/v-s-vernon-jones/text

[2] https://www.deepl.com/translator

to use a single source text and generate the other versions through translation from that source.

The process involved a subsequent crucial step of manual validation and post-editing by native linguist annotators. This manual validation was essential to correct any translation errors, ensure naturalness in the target languages, and, most importantly, guarantee that the semantic and discourse properties central to our annotation were preserved across all three language versions. For the present work, we have just used the Portuguese data to experiment with the newly developed annotation scheme.

3.2. Dataset description

The final corpus is a multilingual parallel collection of 45 Aesop's fables, with each fable available in all three languages: English, European Portuguese, and Polish. The fables were specifically selected for their narrative nature and manageable length, making them suitable for the detailed, manual annotation of fine-grained discourse and temporal features. The consistent number of texts per language provides a balanced foundation for a future comparative study. Table 1 specifies the word count for each fable and the total corpus size.

Fable title	EN	PL	PT
The Ant and the Grasshopper	126	96	141
The Crow and the Pitcher	86	85	85
The Dog in the Manger	64	98	68
The Dog and Its Reflection	86	122	98
The Dog and the Wolf	162	297	175
The Fox and the Grapes	80	63	78
The Frogs Who Desired a King	193	150	180
The Goose that Laid the Golden Eggs	100	103	104
The Lion and the Mouse	175	159	180
The Lion's Share	181	146	164
The Mice and the Weasels	167	149	140
The Tortoise and the Hare	142	145	121
Town Mouse and Country Mouse	259	269	221
The North Wind and the Sun	157	166	121
The Wolf in Sheep's Clothing	96	88	85
Total	2074	2136	2061

Table 1: Word count for various Aesop's fables in different languages

4. The annotation

4.1. The annotation scheme

The Semantic Annotation Framework (SemAF), ISO 24617, comprises multiple components designed to form a comprehensive, unified annotation scheme. However, these components were developed independently and at different times. Although ISO 24617-6 (ISO, 2016a) provides general principles and guidelines for their integration, seeking to address overlaps and inconsistencies across the SemAF components, its scope remains limited. Consequently, when combining different parts of SemAF, as in the present work, inconsistencies and unresolved gaps can be expected and require coherent solutions.

The present annotation scheme builds on several of our previous works, particularly those related to the Text2Story annotation framework (Silvano et al., 2021; Leal et al., 2022; Silvano et al., 2024) and to discourse annotation (Silvano et al., 2023; Silvano and Damova, 2023; Tomaszewska et al., 2024; Silvano et al., 2025). However, as this is the first study in which we attempt to combine discourse and temporal annotation within the ISO 24617 framework, a number of additional procedures had to be followed.

The first step in designing the annotation scheme involved developing strategies to harmonize the two relevant ISO standards, ISO 24617-1 (ISO, 2012) and ISO 24617-8 (ISO, 2016b). One of the main inconsistencies encountered concerned the differing units of annotation: at the discourse layer, arguments are assigned specific roles according to the discourse relation, while at the temporal layer, the annotation is centered on events. To reconcile these perspectives, we introduced a new entity structure, *situation*, corresponding to the discourse arguments, within which we annotated the core temporal, morphosyntactic, and semantic features of the associated events.

The overall architecture of our annotation scheme follows the modular framework of ISO 24617 and comprises two main structural types: entity structures and link structures, each defined by a set of attributes and values, as illustrated in Figure 1.

The annotation of entities encompasses three main categories: *Events*, *Situations*, and *Signals*. In ISO 24617-1 (ISO, 2012), the concept of an *event* corresponds to the notion of an *eventuality* (Bach, 1986), defined as "something that can be said to obtain or hold true, happen, or occur." The markables of *Events* in ISO 24617-1 are typically verbs (e.g., "to dance"), although other word classes, such as nouns, can also express eventualities (e.g., "party").

To characterize different types of eventualities, ISO 24617-1 proposes a set of attributes. For this project, we restricted the annotation to those attributes most relevant for the computation of discourse relations: *Type*, *Tense*, and *Polarity*. The attribute *Type* relates to aspect and follows three values that broadly correspond to the classical classification proposed by Vendler (1957): *state* (a situation in which something obtains or holds true), *process* (a dynamic, atelic situation with duration), and *transition* (a dynamic situation with duration and a consequent state). Events are also anno-

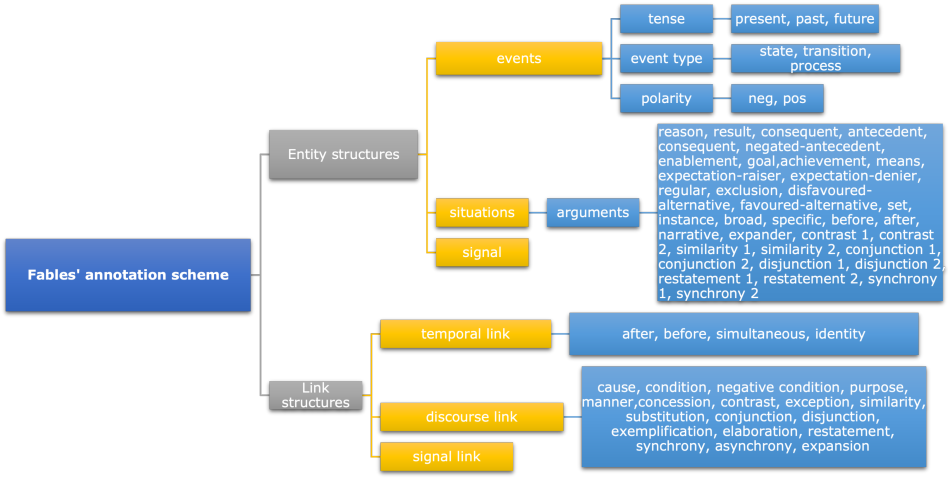

Figure 1: Fables-DTR annotation scheme

tated for *Tense*. In this project, we simplified the ISO 24617-1 specification to include only three values, *past*, *present*, and *future*, to represent the time interval in which the situations are located. Finally, events are annotated for *Polarity*, with two possible values: *positive* and *negative*, indicating whether the situation is asserted or negated, respectively. For example, in the sentence *A wolf decided to disguise itself to attack a flock of sheep*, the event *decided* was annotated with the values *transition, past, positive*.

Situations correspond to *entity structures* that represent discourse arguments. In terms of their *markables*, each situation encompasses the complete syntactic unit associated with the discourse argument, excluding the explicit signal that marks the discourse relation. For example, if the discourse argument is a clause, its markable includes all arguments and adjuncts belonging to that clause, but excludes the preposition, conjunction, or connector introducing it. Using again the *wolf* example, the annotator identifies two situations: "A wolf decided to disguise itself" and "attack a flock of sheep." Regarding their attributes, situations are annotated according to their role within a given discourse relation. The roles used in this project correspond to the DR-core relations defined in ISO 24617-8 (ISO, 2016b). In the *wolf* example, the two situations function as arguments of a *Purpose* discourse relation; therefore, the first situation is annotated as *enablement* and the second as *goal*.

Signals represent the discourse markers (e.g., "because", "although") that explicitly indicate the discourse relation. In our running example, the signal corresponds to *to*.

Our annotation scheme distinguishes two types of *link structures*: (1) *semantic links*, which express temporal and discourse relations between situations; and (2) *technical links*, which serve purely operational purposes related to the annotation platform.

The *semantic links* encode the relationships between situations. Temporal relations follow a subset of those proposed in ISO 24617-1 (ISO, 2012), including *after* and *before* (temporal succession), *simultaneous* (temporal overlap, encompassing partial or total simultaneity as well as inclusion), and *identity* (for situation coreference). Discourse relations are drawn from the DR-core relations specified in ISO 24617-8 (ISO, 2016b). Both discourse and temporal links are always established from the last argument to the preceding one in the linear order of the discourse. In the running example, the second situation is connected to the first by the temporal link *After* and the discourse link *Purpose*.

Finally, two *technical* (non-semantic) links are employed to ensure correct information extraction: the *signal* link connects a signal to the situation it introduces, and the *nucleus-of* link connects an event to the situation it projects. Returning to the *wolf* example, there is a *signal* link between "to" and "attack a flock of sheep", and two *nucleus-of* links: one between the event "decided" and the situation "A wolf decided to disguise itself", and another between the event "attack" and the situation "attack a flock of sheep".

Figure 2 presents our running example fully annotated in the INCEpTION platform (Klie et al., 2018), which was used as the annotation tool in this work.

4.2. The annotation process

The annotation was conducted by two annotators with extensive experience in semantic annotation,

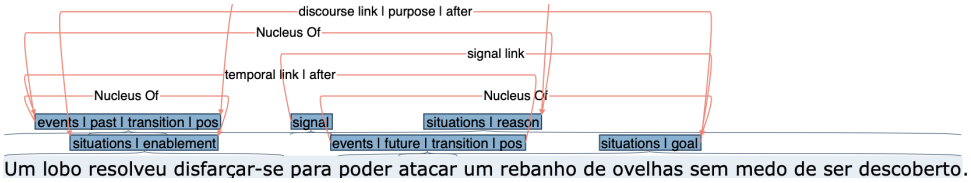

Figure 2: *A wolf decided to disguise himself so that he could attack a flock of sheep without fear of being discovered.* The wolf-example fully annotated in INCEpTION from the Portuguese translated fable "The Wolf in Sheep's Clothing".

particularly in discourse relations using the ISO 24617-8 framework. Before the start of the annotation campaign, several meetings were held to clarify questions regarding specific aspects of the annotation manual and the implementation of the scheme on the chosen platform, INCEpTION (Klie et al., 2018).

Following this preparatory phase, each annotator independently annotated the entire set of fables, adhering to the methodology described in the manual. The process began with the identification and classification of discourse arguments (situations) and their corresponding signals (when present). Annotators then established the discourse and signal links. In a subsequent phase, they identified the event corresponding to each situation, annotated its attributes, and linked it to the relevant situation. Finally, the annotator established the temporal link between the events included in each situation connected by a discourse relation. After completion, all annotations underwent curation by a researcher with extensive expertise in semantic annotation, especially in tense and discourse relations. The preliminary curation results were analyzed to detect and correct any annotation errors or inconsistencies. Once these revisions were completed, the final curated dataset was produced, the results of which are presented in this paper.

For the present study, we used only the Portuguese data to test and validate the newly developed annotation scheme.

To facilitate reproducibility, the texts in the three languages, the annotated Portuguese dataset, and the annotation guidelines are made available through a GitHub repository[3]. The resources are released under the Apache License, Version 2.0, a permissive license that allows the use of the annotations, including for commercial purposes.

4.3. Inter-annotator agreement

Inter-annotator agreement was calculated using Krippendorff's alpha, a standard metric suitable for datasets with class imbalance (Artstein and Poesio, 2008). Agreement scores between the two annotators of the Portuguese Aesop's Fables corpus (ANN1 and ANN2), before curation are presented in Figure 3. To compute agreement, we considered span matches that allowed at most one word of difference, preserving the word sequence. Specifically, we applied the Longest Common Subsequence (LCS) algorithm. For link annotations, both source and target spans were allowed a one-word difference to qualify as a match.

Figure 3 shows that the IAA is relatively low for both discourse and temporal links, although it is slightly higher for the temporal layer than for the discourse relations layer. The low IAA can be partly explained by the density and complexity of the annotation task. In this scheme, annotators are required to perform an exhaustive annotation of all discourse relations in the fables together with their associated temporal relations. As a result, the same text segment may simultaneously function as an argument of a discourse relation while also being part of a larger argument involved in another relation. Previous work focusing exclusively on temporal relations (Fernandes et al., 2025) or exclusively on discourse relations (Tomaszewska et al., 2024) grounded in ISO 24617-1 (ISO, 2012) and ISO 24617-8 (ISO, 2016b) reports substantially higher agreement scores. This observation suggests that the lower IAA obtained in our study may stem from the increased difficulty of the proposed scheme. By combining discourse and temporal relations within the same annotation framework, the scheme introduces a more intricate annotation structure, which in turn raises the level of difficulty of the task for annotators. Additionally, the results may also be explained by the fact that, as noted in (Tomaszewska et al., 2024), the set of discourse relations proposed in ISO 24617-8 sometimes lacks sufficient specificity, particularly with regard to distinctions between certain relations, such as ELABORATION and EXPANSION. This limited granularity may lead annotators to adopt different interpretations when selecting the most appropriate relation, thereby contributing to lower agreement scores. In the future, we will overpass this problems by adding to the guidelines more specifications to decide between discourse and temporal relations.

An example illustrating disagreement regarding

[3] https://github.com/LIAAD/fables

both DLinks and TLinks is the following: *. . . quando o rebanho foi recolhido para passar a noite, ele foi fechado com as outras ovelhas. . .* ('...when the flock was gathered for the night, he was shut in with the other sheep. . . '). In this case, ANN1 considered the discourse relation between *when the flock was gathered for the night* and *he was shut in with the other sheep* to be Synchrony (the situations overlap each other), whereas ANN2 classified it as Asynchrony (one situation precedes the other temporally). Since these temporal discourse relations are associated with specific temporal relations, namely simultaneity and temporal succession, respectively, this difference in interpretation also resulted in disagreement in the corresponding TLinks. A plausible explanation for this divergence lies in the temporal semantics of the clause introduced by *quando*. One annotator may have interpreted the gathering of the flock and the act of shutting the wolf in as part of the same temporal frame, thus signaling simultaneity, while the other may have interpreted the first event as temporally preceding and enabling the second, leading to an interpretation of succession. In the first case, the annotator appears to focus on the resultant state of the flock being gathered, whereas the other annotator interprets the same situation with a more eventive reading.

In another example, *"os ratos levavam sempre a pior, sendo muitos deles mortos"* ('the mice always came off worse, many of them being killed'), ANN1 classified the discourse relation as ELABORATION, whereas ANN2 annotated it as CAUSE. A plausible explanation for this disagreement lies in the interpretation of the participial clause *sendo muitos deles mortos*. One annotator may have interpreted this clause as providing additional information that further specifies the previous statement that the mice were always worse off, thus functioning as an elaboration. The other annotator may instead have interpreted the clause as expressing the reason why the mice consistently fared worse, namely that many of them were killed, leading to a causal interpretation.

Regarding the event-related attributes, the agreement was moderate, which is expected, as these tasks are generally more objective and classificatory in nature. In contrast, the identification of markables corresponding to situations, signals, and events proved more challenging, with the signal category showing the lowest agreement. This may be explained by some ambiguity in the annotation manual regarding which linguistic forms and expressions should be marked as signals. To address this issue, additional illustrative examples will be included in the next version of the manual. Nonetheless, the curator resolved these discrepancies consistently across the dataset.

5. Fables-DTR Dataset characterization

5.1. Overall analysis

A total of 210 discourse relations (DRs) were annotated, of which 135 included an explicit signal (64.3%) and 75 did not (35.7%), indicating a predominance of explicit DRs. These relations involved 228 situations and 228 events, as summarized in Table 2.

Table 2: Frequency of DRs, temporal relations, signals, and events.

Category	Count
DR	210
DR with signal	135
DR without signal	75
Temporal Relations	210
Signals	135
Events	228

Regarding the most frequent DR types, Table 3 shows that ASYNCHRONY is by far the most common relation, which is consistent with the narrative nature of the fables. EXPANSION and CAUSE follow in frequency. The former is typically used to provide additional characterization of a story element (such as a character or situation), and the latter to explain characters' actions or motivations.

Table 3: Frequency of DRs with and without signal.

DR	With signal	Without signal	Total
Asynchrony	62	34	96
Expansion	14	11	25
Cause	17	10	27
Concession	15	1	16
Purpose	15	1	16
Synchrony	6	8	14
Elaboration	3	8	11
Manner	1	2	3
Conjunction	2	0	2

With respect to the events' attributes, almost all events correspond to positive predications (218 vs. 10), and the predominant aspectual class is Transitions (139), although there is a reasonable number of Processes (43) and States (46).

5.2. Detailed analysis

In this section, we aim to demonstrate how the performed annotation contributes to a finer characterization of discourse relations in this type of text, fables, by identifying the most relevant features of the discourse relations and their arguments. We begin this analysis by focusing on the signals and temporal relations systematized in Table 4.

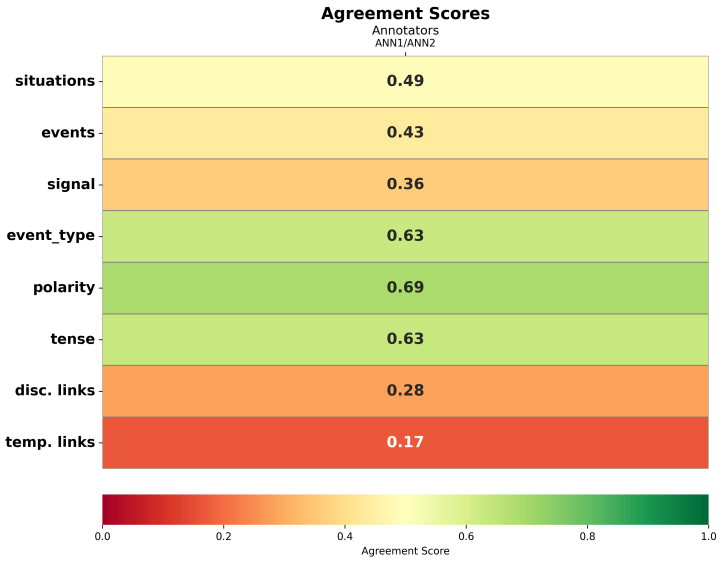

Figure 3: Inter-annotator agreement of Entities, Attributes and Links in the Portuguese data

Table 4: Frequency of temporal relations for each DR. Abbreviations: w/ sig. = with signal, w/o sig. = without signal, Sim. = Simultaneous, Id. = Identity.

DR	Signal	After	Sim.	Before	Id.
Asynchrony	w/ sig.	62			
	w/o sig.	34			
Synchrony	w/ sig.		6		
	w/o sig.		8		
Expansion	w/ sig.	5	7	2	
	w/o sig.	2	7	2	
Concession	w/ sig.	12	3		
	w/o sig.		1		
Purpose	w/ sig.	13	1	1	
	w/o sig.	1			
Cause	w/ sig.	7	8	2	
	w/o sig.	4	6		
Conjunction	w/ sig.		2		
	w/o sig.				
Elaboration	w/ sig.		3		
	w/o sig.		2		6
Manner	w/ sig.		1		
	w/o sig.		2		

The results show that, irrespective of whether a signal is present, ASYNCHRONY is systematically associated with temporal succession *After*, while SYNCHRONY is consistently associated with *Simultaneous*. This association was expected, since these DRs inherently imply these temporal relations, regardless of the presence or absence of a signal. A distinct pattern emerges with EXPANSION, where the temporal relation tends to be *Simultaneous*, although all temporal relations are attested (again, independently of signaling). CONCESSION and PURPOSE exhibit similar distributions: both are predominantly associated with *After*, suggesting that these DRs are primarily linked to temporal succession. Moreover, they are typically signal-marked, thus occurring mainly as explicit DRs. CAUSE also tends to be explicit, although to a lesser degree than CONCESSION and PURPOSE. It differs from these DRs in that it shows a more balanced distribution between the temporal relations *Simultaneous* and *After*, with a slight predominance of the former. Finally, the DRs CONJUNCTION, ELABORATION, and MANNER, which are less frequent overall, are exclusively associated with the temporal relation *Simultaneous*. It is worth recalling that, in the case of ELABORATION, the six occurrences annotated as *Identity* correspond to two descriptions of the same situation. This implies that both descriptions refer to a single relevant temporal interval, bringing the relation of *Identity* conceptually close to that of *Simultaneous*.

Aspectual classes constitute another key component of our annotation scheme. Table 5 presents the results concerning the aspectual class of situations, their role as arguments of discourse relations (DRs), and the presence or absence of a signal.

The data show that, for certain DRs, the presence or absence of a signal does not affect the distribution of aspectual classes. This is the case with ASYNCHRONY, where *Transitions* predominate in both DR arguments, and with SYNCHRONY, where the three aspectual classes are more evenly distributed across both arguments. By contrast, in EXPANSION, the existence or absence of a signal appears to influence the distribution of aspectual classes. Although some dispersion is observed, situations tend to be *Transitions* (i.e., telic situations) when a signal is present, whereas they tend to be *Processes* and *States* (atelic situations) when no signal is present. CONCESSION and PURPOSE once again display similar tendencies: situations tend to be *Transitions* in both arguments. The same pattern is observed for ELABORATION. The DR CAUSE

Table 5: Frequency of combinations of event attributes DR (with arguments) / aspectual class. Abbreviations: trans. = transition, proc. = process.

DR	Argument	With signal			Without signal		
		Trans.	Proc.	State	Trans.	Proc.	State
Asynchrony	After	44 (70.97%)	13 (20.97%)	5 (8.06%)	30 (88.24%)	3 (8.82%)	1 (2.94%)
	Before	47 (75.81%)	10 (16.13%)	5 (8.06%)	29 (85.29%)	3 (8.82%)	2 (5.88%)
Synchrony	Synchrony 1	3 (50.00%)	1 (16.67%)	2 (33.33%)	2 (25.00%)	2 (25.00%)	4 (50.00%)
	Synchrony 2	3 (50.00%)	2 (33.33%)	1 (16.67%)	3 (37.50%)	3 (37.50%)	2 (25.00%)
Expansion	Expander	9 (64.29%)	3 (21.43%)	2 (14.29%)	3 (27.27%)	3 (27.27%)	5 (45.45%)
	Narrative	8 (57.14%)	2 (14.29%)	4 (28.57%)	5 (45.45%)		6 (54.55%)
Concession	Expectation-denier	11 (73.33%)	1 (6.67%)	3 (20.00%)			1 (100.00%)
	Expectation-raiser	10 (66.67%)	4 (26.67%)	1 (6.67%)		1 (100.00%)	
Purpose	Enablement	12 (80.00%)	2 (13.33%)	1 (6.67%)	1 (100.00%)		
	Goal	9 (60.00%)	5 (33.33%)	1 (6.67%)			1 (100.00%)
Cause	Reason	6 (35.29%)	2 (11.76%)	9 (52.94%)	5 (50.00%)		5 (50.00%)
	Result	7 (41.18%)	1 (5.88%)	9 (52.94%)	9 (90.00%)		1 (10.00%)
Conjunction	Conjunction 1	1 (50.00%)		1 (50.00%)			
	Conjunction 2		2 (100.00%)				
Elaboration	Broad	3 (100.00%)			7 (87.50%)	1 (12.50%)	
	Specific	1 (33.33%)	2 (66.67%)		5 (62.50%)	2 (25.00%)	1 (12.50%)
Manner	Achievement	1 (100.00%)				2 (100.00%)	
	Means			1 (100.00%)		2 (100.00%)	

also reveals a certain asymmetry linking aspectual classes to the presence or absence of a signal. With a signal, CAUSE arguments exhibit a slightly higher proportion of *States* than *Transitions*; without a signal, the inverse pattern is found, with *Transitions* slightly predominating. Finally, the DRs CONJUNCTION and MANNER occur too infrequently to allow for the identification of robust trends.

The analysis of event attributes in conjunction with TENSE and POLARITY did not yield any particularly significant results. Regarding TENSE, most situations occur in the *Past*, which limits the interpretability of this feature. Instances annotated as *None* correspond to nominal events or verbs in non-finite forms. However, these are relatively few in number (except in the case of PURPOSE–GOAL with a signal). Only a single occurrence of the *Future* tense was identified. As for POLARITY, the *Positive* value predominates overall. A slightly higher number of *Negative* instances is observed in DRs with a signal than in those without. These negative cases are distributed across different DRs and typically occur in both DR arguments. Tables 8 and 9 in the summarize these results.

Finally, we examined the distribution of discourse signals, as summarized in Table 6. Despite the wide variety of signal types observed (prepositions, conjunctions, and interphrasal connectors), the results show a clear concentration on five forms: *e* ('and') (copulative conjunction), *para* ('to/for') (preposition), *mas* ('but') (adversative conjunction), *que* ('who/that') (relative pronoun or conjunction), and *quando* ('when') (temporal conjunction).

Table 6: Frequency of Signal markables

Markable	Count	Markable	Count
e (and)	50	mal (as soon as)	1
para (in order to)	15	nem (nor)	1
mas (but)	14	então (then)	1
que (that)	13	no qual (in which)	1
quando (when)	12	em seguida (next)	1
por (because of)	4	tanto que (so much that)	1
por fim (finally)	3	entretanto (meanwhile)	1
depois (afterwards)	3	como se (as if)	1
por isso (therefore)	2	primeiro (first)	1
assim (thus)	2	porque (because)	1
antes de (before)	2	pois (since)	1
embora (although)	1	pelo que (so that)	1
até que (until)	1	onde (where)	1

TOTAL = 135

Signals typically occur by marking the same argument of the discourse relation (DR), as shown in Table 7. This pattern is observed for several DRs: ASYNCHRONY (*After*), EXPANSION (*Expander*), SYNCHRONY (*Synchrony 2*), CONCESSION (*Expectation-*

raiser), PURPOSE (*Goal*), and CONJUNCTION (*Conjunction 2*). This generalization is less consistent for the DR CAUSE, where the signal is associated with the *Reason* argument in six instances and with the *Result* argument in eleven.

Table 7: Signal distribution by DR and argument.

DR	Argument	Count
Asynchrony	After	52
	Before	10
Expansion	Expander	13
	Narrative	1
Synchrony	Synchrony 2	5
	Synchrony 1	1
Concession	Expectation-denier	14
	Expectation-raiser	1
Purpose	Goal	15
Cause	Reason	6
	Result	11
Conjunction	Conjunction 2	2
Elaboration	Broad	1
	Specific	2
Manner	Means	1

Analyzing the distribution of signals across DR arguments in Table 10, we observe that, in the case of ASYNCHRONY, despite a considerable variety of signals, the conjunction *e* ('and') is clearly preferred to mark the *After* argument, while there is a tendency for the conjunction *quando* ('when') to mark the *Before* argument. The DR EXPANSION also presents a range of signals, essentially *wh* words, although to a lesser extent than ASYNCHRONY. The *Expander* argument tends to be marked by *que* ('who/that'). For CAUSE, the *Reason* argument is most frequently marked by the preposition *por* ('by'/'because of'), while the *Result* argument shows more variation, with a slight preference for *que* ('that'). CONCESSION and PURPOSE stand out for systematically marking the same argument with the same signal: in CONCESSION, the conjunction *mas* ('but') predominates, whereas in PURPOSE the preposition *para* ('to/for') is consistently used. As for the remaining DRs, their low frequency prevents the identification of reliable tendencies.

6. Concluding remarks

The FABLES-DTR corpus bridges a significant gap in existing discourse and temporal annotation resources by integrating ISO 24617-1 (Temporal Relations) and ISO 24617-8 (Discourse Relations) within a single interoperable framework. Unlike earlier corpora that annotated discourse or temporal layers separately, FABLES-DTR systematically links discourse relations, such as PURPOSE, CAUSE, CONCESSION, and ASYNCHRONY, to their corresponding temporal profiles and aspectual attributes.

The quantitative analysis of the Portuguese fables subset reveals that explicit discourse relations are the most frequent, with ASYNCHRONY, EXPANSION, and CAUSE emerging as the most recurrent types. Within ASYNCHRONY, temporal succession relations play a central role in establishing narrative coherence, reflecting the intrinsic sequentiality of fable narratives. As for aspectual features, *Transitions* predominate across DRs. Explicit connectives occur most often in PURPOSE and CONCESSION relations, whereas implicit DRs are more frequent in ASYNCHRONY and ELABORATION.

Ultimately, FABLES-DTR demonstrates that the integration of discourse and temporal annotation within a unified framework is both feasible and enables relevant conclusions about discourse relations. This resource provides valuable insights into the interplay between discourse structure and temporal sequencing, and offers a robust empirical foundation for advancing multilingual discourse parsing, narrative temporal modeling, and the automatic identification of implicit relations. Moreover, it establishes a strong benchmark for future ISO-compliant annotation initiatives.

7. Acknowledgements

This research was funded in part by the National Science Centre, Poland, grant 2023/50/A/HS2/00559 ("Universal Discourse: a multilingual model of discourse relations"). This work is also financed by National Funds through the FCT - Fundação para a Ciência e a Tecnologia, I.P. (Portuguese Foundation for Science and Technology) within the project StorySense, with reference 2022.09312.PTDC (DOI al10.54499/2022.09312.PTDC).

8. Bibliographical References

Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational linguistics*, 34(4):555–596.

Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, United States.

Emmon Bach. 1986. The algebra of events. *Linguistics and philosophy*, pages 5–16.

Johan Bos, Valerio Basile, Kees Evang, Nissim J. Venhuizen, and Johannes Bjerva. 2017. The groningen meaning bank. In Nancy Ide and James Pustejovsky, editors, *Handbook of Linguistic Annotation*, pages 463–496. Springer, Dordrecht.

Tommaso Caselli and Piek Vossen. 2017. The event StoryLine corpus: A new benchmark for causal and temporal relation extraction. In *Proceedings of the Events and Stories in the News Workshop*, pages 77–86, Vancouver, Canada. Association for Computational Linguistics.

Ana Luísa Fernandes, Purificação Silvano, António Leal, Nuno Guimarães, Rita Rb-Silva, Luís Filipe Cunha, and Alípio Jorge. 2025. The incremental process of building an annotation scheme for clinical narratives in Portuguese: the contribution of human variation analysis. In *Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 332–343, Vienna, Austria. Association for Computational Linguistics.

ISO. 2012. Language resource management—semantic annotation framework (semaf) — part 1: Time and events (semaf-time, iso-timeml). ISO Standard ISO 24617-1:2012. Https://www.iso.org/standard/37326.html.

ISO. 2016a. Language resource management — semantic annotation framework — part 6: Principles of semantic annotation (semaf principles). ISO Standard ISO 24617-6:2016. Available at: https://www.iso.org/standard/37327.html.

ISO. 2016b. ISO 24617-8. 2016. Language resource management, part 8: Semantic relations in discourse (DR-Core). Standard, International Organization for Standardization, Geneva, CH.

Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic: Introduction to Model-Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, volume 42 of *Studies in Linguistics and Philosophy*. Kluwer Academic Publishers, Dordrecht.

Kimi Kaneko, Saku Sugawara, Koji Mineshima, and Daisuke Bekki. 2016. Annotation and analysis of discourse relations, temporal relations and multi-layered situational relations in Japanese texts. In *Proceedings of the 12th Workshop on Asian Language Resources (ALR12)*, pages 10–19, Osaka, Japan. The COLING 2016 Organizing Committee.

Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart De Castilho, and Iryna Gurevych. 2018. The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th international conference on computational linguistics: system demonstrations*, pages 5–9.

António Leal, Purificação Silvano, Evelin Amorim, Inês Cantante, Fátima Silva, Alípio Mario Jorge, and Ricardo Campos. 2022. The place of ISO-space in Text2Story multilayer annotation scheme. In *Proceedings of the 18th Joint ACL - ISO Workshop on Interoperable Semantic Annotation within LREC2022*, pages 61–70, Marseille, France. European Language Resources Association.

William Mann and Sandra Thompson. 1988. Rethorical structure theory: Toward a functional theory of text organization. *Text*, 8:243–281.

Paramita Mirza, Rachele Sprugnoli, Sara Tonelli, and Manuela Speranza. 2014. Annotating causality in the TempEval-3 corpus. In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)*, pages 10–19, Gothenburg, Sweden. Association for Computational Linguistics.

Maciej Ogrodniczuk, Aleksandra Tomaszewska, Daniel Ziembicki, Sebastian Żurowski, Ryszard Tuora, and Aleksandra Zwierzchowska. forthcoming. Polish Discourse Corpus (PDC): Corpus Design, ISO-Compliant Annotation, Data Highlights, and Parser Development. In *Proceedings of The 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC COLING 2024)*.

James Pustejovsky, Marc Verhagen, Roser Saurí, Jessica Littman, Robert Gaizauskas, Graham Katz, Inderjeet Mani, Robert Knippen, and Andrea Setzer. 2006. Timebank 1.2. Linguistic Data Consortium, LDC2006T08. Release Date: April 17, 2006.

Purificação Silvano, Evelin Amorim, António Leal, Inês Cantante, Alípio Jorge, Ricardo Campos, and Nana Yu. 2024. Untangling a web of temporal relations in news articles. In *Proceedings of the Text2Story'24 Workshop*. Glasgow, Scotland, 24 March 2024 (workshop proceedings).

Purificação Silvano, João Cordeiro, António Leal, and Sebastião Pais. 2023. DRIPPS: a corpus with discourse relations in perfect participial sentences. In *Proceedings of the 4th Conference on Language, Data and Knowledge*, pages 470–481, Vienna, Austria. NOVA CLUNL, Portugal.

Purificação Silvano and Mariana Damova. 2023. ISO-DR-core plugs into ISO-dialogue acts for a cross-linguistic taxonomy of discourse markers. In *Proceedings of the 4th Conference on Language, Data and Knowledge*, pages 440–448, Vienna, Austria. NOVA CLUNL, Portugal.

Purificação Silvano, António Leal, Fátima Silva, Inês Cantante, Fatima Oliveira, and Alípio

Mario Jorge. 2021. Developing a multilayer semantic annotation scheme based on ISO standards for the visualization of a newswire corpus. In *Proceedings of the 17th Joint ACL - ISO Workshop on Interoperable Semantic Annotation*, pages 1–13, Groningen, The Netherlands (online). Association for Computational Linguistics.

Purificação Silvano, Giedrė Valūnaitė Oleškevičienė, Chaya Liebeskind, and Mariana Damova. 2025. Questions in monologues: An analysis grounded on iso. *Linguistics Vanguard*.

Aleksandra Tomaszewska, Purificação Silvano, António Leal, and Evelin Amorim. 2024. ISO 24617-8 applied: Insights from multilingual discourse relations annotation in English, Polish, and Portuguese. In *Proceedings of the 20th Joint ACL - ISO Workshop on Interoperable Semantic Annotation @ LREC-COLING 2024*, pages 99–110, Torino, Italia. ELRA and ICCL.

Zeno Vendler. 1957. Verbs and times. *The philosophical review*, 66(2):143–160.

9. Appendices

Table 8: Tense distribution by DR argument, with and without signal.

DR	Argument	With signal			Without signal		
		Past	Future	None	Past	Future	None
Asynchrony	After	59 (95.16%)		3 (4.84%)	34 (100.00%)		
	Before	59 (95.16%)		3 (4.84%)	30 (88.24%)		4 (11.76%)
Synchrony	Synchrony 1	6 (100.00%)			7 (87.50%)		1 (12.50%)
	Synchrony 2	6 (100.00%)			5 (62.50%)		3 (37.50%)
Expansion	Expander	14 (100.00%)			10 (90.91%)		1 (9.09%)
	Narrative	13 (92.86%)		1 (7.14%)	10 (90.91%)		1 (9.09%)
Concession	Expectation-denier	13 (86.67%)		2 (13.33%)	1 (100.00%)		
	Expectation-raiser	15 (100.00%)			1 (100.00%)		
Purpose	Enablement	13 (86.67%)		2 (13.33%)	1 (100.00%)		
	Goal		1 (6.67%)	14 (93.33%)			1 (100.00%)
Cause	Reason	13 (76.47%)		4 (23.53%)	3 (30.00%)		7 (70.00%)
	Result	16 (94.12%)		1 (5.88%)	10 (100.00%)		
Conjunction	Conjunction 1	2 (100.00%)					
	Conjunction 2	2 (100.00%)					
Elaboration	Broad	3 (100.00%)			7 (87.50%)		1 (12.50%)
	Specific	3 (100.00%)			6 (75.00%)		2 (25.00%)
Manner	Achievement	1 (100.00%)			2 (100.00%)		
	Means	1 (100.00%)					2 (100.00%)

Table 9: Polarity distribution by DR argument, with and without signal.

DR	Argument	With signal		Without signal	
		pos	neg	pos	neg
Asynchrony	After	61 (98.39%)	1 (1.61%)	34 (100.00%)	
	Before	60 (96.77%)	2 (3.23%)	33 (97.06%)	1 (2.94%)
Expansion	Expander	14 (100.00%)		10 (90.91%)	1 (9.09%)
	Narrative	13 (92.86%)	1 (7.14%)	11 (100.00%)	
Synchrony	Synchrony 1	6 (100.00%)		8 (100.00%)	
	Synchrony 2	6 (100.00%)		7 (87.50%)	1 (12.50%)
Concession	Expectation-denier	14 (93.33%)	1 (6.67%)		1 (100.00%)
	Expectation-raiser	15 (100.00%)		1 (100.00%)	
Purpose	Enablement	15 (100.00%)		1 (100.00%)	
	Goal	15 (100.00%)		1 (100.00%)	
Cause	Reason	16 (94.12%)	1 (5.88%)	10 (100.00%)	
	Result	13 (76.47%)	4 (23.53%)	10 (100.00%)	
Conjunction	Conjunction 1	1 (50.00%)	1 (50.00%)		
	Conjunction 2	1 (50.00%)	1 (50.00%)		
Elaboration	Broad	3 (100.00%)		8 (100.00%)	
	Specific	3 (100.00%)		8 (100.00%)	
Manner	Achievement	1 (100.00%)		2 (100.00%)	
	Means	1 (100.00%)		2 (100.00%)	

Table 10: Signal markables frequency used per DR and argument.

DR	Argument	Markable	Count
Asynchrony	After	e (and)	43
	After	por fim (finally)	3
	After	até que (until)	1
	After	depois (afterwards)	3
	After	em seguida (next)	1
	After	assim (thus)	1
	Before	quando (when)	6
	Before	e (and)	1
	Before	mal (as soon as)	1
	Before	antes de (before)	2
Expansion	Expander	e (and)	1
	Expander	que (that)	9
	Expander	onde (where)	1
	Expander	então (then)	1
	Expander	no qual (in which)	1
	Narrative	quando (when)	1
Synchrony	Synchrony 2	quando (when)	3
	Synchrony 2	e (and)	1
	Synchrony 2	entretanto (however/meanwhile)	1
	Synchrony 1	quando (when)	1
Concession	Expectation-denier	mas (but)	14
	Expectation-raiser	embora (although)	1
Purpose	Goal	para (in order to)	15
Cause	Result	que (that)	4
	Result	por isso (therefore)	2
	Result	pelo que (therefore/so that)	1
	Result	e (and)	2
	Result	assim (thus)	1
	Result	tanto que (so much that)	1
	Reason	porque (because)	1
	Reason	pois (since)	1
	Reason	por (because of)	4
Conjunction	Conjunction 2	e (and)	1
	Conjunction 2	nem (nor)	1
Elaboration	Broad	quando (when)	1
	Specific	e (and)	1
	Specific	primeiro (first)	1
Manner	Means	como se (as if)	1