

# CATS: An annotation scheme of causality and temporal structure

Nana Yu\*, Purificação Silvano\*<sup>†§</sup>, Luís Filipe Cunha<sup>†§</sup>, Alípio Jorge<sup>†§</sup>

\*Faculty of Arts and Humanities, University of Porto

† Faculty of Sciences, University of Porto

‡Centre for Linguistics of the University of Porto (CLUP)

§ INESC TEC

Porto, Portugal

robertananayu@hotmail.com, {purificacao.silvano, luis.f.cunha, amjorge}@inesctec.pt

## Abstract

This paper presents CATS, a causal and temporal annotation scheme designed to jointly represent causal relations and temporal structures in news texts. The proposed framework integrates components of ISO 24617 Semantic Annotation Framework (SemAF), drawing in particular on Part 1 (Time and Events) (ISO 24617-1: 2012) and Part 8 (Semantic Relations in Discourse) (ISO 24617-8: 2016). Building on the Text2Story annotation framework (Silvano et al., 2021), the scheme adapts and extends its principles for representing temporal information while introducing new entities and links for modeling causal relations. The resulting annotation model enables the integrated representation of causal arguments, events, temporal relations, and causal signals within a unified structure. By jointly capturing causal and temporal dependencies, CATS provides a resource for studying the interaction between causality and temporality in discourse and supports downstream NLP tasks such as event extraction, temporal ordering, and causal reasoning.

**Keywords:** causal relations, temporal relations, discourse relations, semantic annotation schemes, ISO 24617

## 1. Introduction

Causality constitutes one of the fundamental principles guiding human reasoning, as knowledge about the world tends to be organized through causal networks that explain, justify, or constrain the occurrence of phenomena (Pearl, 2009; Woodward, 2003; Sloman, 2005). For this reason, causality has been the subject of study across several areas of social sciences, humanities, and natural sciences, and, more recently, has also become an important topic in the field of Natural Language Processing (NLP) (Yang et al., 2022; Feder et al., 2022; Koupae et al., 2025).

Language widely reflects causality relations, but linguistic studies acknowledge that defining causality is a complex task (Talmy, 2000). In Linguistics, causality has been widely studied within the framework of discourse relations established between textual segments/utterances (Hobbs, 1985; Mann and Thompson, 1988; Asher and Lascarides, 2003; Kehler, 2002). These relations, also referred to as rhetorical relations, constitute an important object of study in semantics, since they explain how different discourse units are connected in the construction of the overall meaning of a text (Mann and Thompson, 1988; Asher and Lascarides, 2003). Several theoretical and annotation frameworks integrate such relations into their discourse models, including proposals by Hobbs (1985), Mann and Thompson (1988), Asher and Lascarides (2003), Prasad et al. (2008), and the ISO discourse-relation

standard ISO 24617-8: 2016 (Prasad and Bunt, 2015; Bunt and Prasad, 2016). Within this perspective, linguistic elements and structures play a central role in inferring discourse relations. In practice, these relations can be signaled or inferred from multiple types of linguistic cues, such as lexical choices (nouns, verbs, adverbs, prepositions), discourse connectives, and syntactic constructions of coordination and subordination (Prasad et al., 2008). Beyond these elements, other linguistic dimensions, such as tense and aspect, also contribute substantially to the interpretation of relations between events (Moens and Steedman, 1988; Lascarides and Asher, 1993). Accordingly, characterizing temporal relations and the aspectual properties of the involved situations is an essential dimension for the analysis of causality. Their interaction has been explicitly discussed in both linguistic theory and NLP-oriented studies of event relations (Lascarides and Asher, 1993; Mirza and Tonelli, 2014; Ning et al., 2018a).

The representation and processing of causal and temporal information have received increasing attention in recent years. Nevertheless, many studies tend to focus primarily on one of these dimensions rather than the other, and the number of works addressing temporal and causal relations jointly remains relatively limited (Bethard et al., 2008; Ning et al., 2018a; Mirza and Tonelli, 2014). Within the domain of semantic annotation schemes, several proposals include labels and attributes designed to represent temporal, aspectual, and discourse-

related properties. In the domain of temporality, relevant work includes Pustejovsky et al. (2003a), Setzer (2001), Mani and Schiffman (2004), Ning et al. (2018b), and ISO 24617-1: 2012 (Pustejovsky et al., 2010), and, more recent ISO-based, multi-layer efforts applied to narrative/news data such as Silvano et al. (2024). In the domain of causality, reference points include the SemEval annotation task that explicitly covers Cause–Effect relations (Girju et al., 2007), corpus-based causal annotation schemes grounded in linguistic analysis (Dunietz et al., 2015, 2017), and recent event-causality resources such as the Causal News Corpus (Tan et al., 2022). Finally, some approaches seek to integrate temporal and causal relations simultaneously, including Bethard et al. (2008) and Mirza and Tonelli (2014).

Despite these contributions, there remains a need for annotation schemes that systematically model how causality and temporality interact in discourse, motivating the development of the proposal presented in this work. The need for a unified treatment of causality and temporality arises from the fact that these two dimensions are deeply interdependent in discourse, yet are typically treated in isolation in existing annotation frameworks. A unified annotation scheme enables the systematic investigation of questions such as how causal relations are temporally anchored, how different types of causal relations correlate with temporal configurations, and how ambiguities between mere temporal succession and causality are resolved in discourse. It also allows the identification of the linguistic and structural cues that signal these interactions, including connectives, tense–aspect marking, and discourse organization. By explicitly representing the interaction between causality and temporality, the present work provides a framework for both fine-grained linguistic analysis of causality and more robust modeling in downstream applications, such as event ordering, causal inference, and narrative understanding.

The main contributions of this paper are the following:

- an integrated scheme for the annotation of causal and temporal relations grounded in ISO 24617;
- a description of the methodology used to develop the scheme;
- a critical analysis of several incompatibilities between the two ISO parts and suggestions for addressing these issues.

The remainder of the paper is organized as follows. Section 2 reviews related work on causal and temporal annotation. Section 3 describes the process of designing the annotation scheme, starting

with the methodology for developing the scheme (Section 3.1), and the motivation for using ISO 24617 (Section 3.2.). Section 3.3 discusses the integration of both layers and presents the main challenges encountered during the process, together with the proposed solutions. Section 4 introduces the final causal and temporal annotation scheme. Finally, Section 5 concludes the paper and outlines directions for future research.

## 2. Related work

A major milestone in the annotation of causal relations in discourse is the work of Prasad et al. (2008) associated with the Penn Discourse Treebank (PDTB). Built from texts from the *Wall Street Journal*, the PDTB includes thousands of annotated instances of causal relations within the semantic class CONTINGENCY. The framework distinguishes between explicit and implicit discourse relations. In explicit relations, causality is marked by lexical connectives (e.g., *because*, *as a result*), which are classified according to their syntactic and functional properties. Annotation follows the syntactic attachment of the connective: the argument syntactically linked to the marker is labeled *Arg2* (typically interpreted as the cause), while the other segment is labeled *Arg1* (often corresponding to the effect). In implicit relations in the PDTB, where no lexical connective is present, annotators identify two arguments and insert an implicit connective between them. *Arg2* corresponds to the clause where the connective would appear, while *Arg1* is the other argument, typically preceding *Arg2* in the text. This approach highlights the discourse-based nature of causal interpretation and the methodological challenges posed by non-lexicalized relations.

An alternative perspective focusing on the linguistic realization of causality is proposed by Dunietz et al. (2015). The authors argue that no single representation scheme can capture the full semantic diversity of causal expressions, since causality is a cognitive construct interacting with dimensions such as temporality, counterfactuality, factuality, and negation. Their annotation framework is grounded in Construction Grammar and defines causal language as any construction that presents one event or state as promoting or hindering another, provided that at least one lexical trigger is present. Consequently, the annotation scheme is restricted to explicit causal relations, excluding implicit cases or constructions whose causal interpretation depends primarily on extralinguistic knowledge. The resulting corpus, *BECauSE 1.0*, identifies three core components, *connective*, *cause*, and *effect*, and distinguishes four causal types: *consequence*, *motivation*, *purpose*, and *inference*.

This line of work is further extended in *BE-*

*CauSE 2.0* (Dunietz et al., 2017), which conceptualizes causality as causal networks in which phenomena may cause, enable, or prevent other events. The updated scheme expands the scope of annotation by capturing overlapping semantic relations frequently associated with causal expressions, including temporal, hypothetical, and contextual relations. Compared to the first version, the scheme reorganizes causal categories, retaining *consequence*, *motivation*, and *purpose* while removing *inference*, and introducing a distinction between facilitating and inhibitory causation. The resulting corpus contains sentences with annotated causal instances and overlapping relations, illustrating the complexity of causal meaning in natural language.

From an event-oriented perspective, Tan et al. (2022) emphasize the central role of causality in natural language understanding and inference tasks. They propose annotation guidelines that classify sentences as either causal or non-causal, allowing both explicit and implicit causal relations as long as both cause and effect are present. Event identification follows the TimeML framework, while the notion of causality is inspired by the CONTINGENCY relation in PDTB 3.0. The resulting *Causal News Corpus* contains news texts annotated with binary sentence-level labels indicating the presence of causal relations.

In contrast to the relatively limited number of causal annotation frameworks, temporal annotation has been extensively studied. A foundational proposal is TimeML (Pustejovsky et al., 2003b), designed to represent events, temporal expressions, and relations between them through a structured set of entities and links. TimeML has been widely adopted and later standardized as ISO-TimeML (ISO 24617-1: 2012). Adaptations for other languages include *TimeBankPT* (Costa and Branco, 2012a,b) for European Portuguese. Building on the ISO Semantic Annotation Framework, the Text2Story scheme (Silvano et al., 2021; Leal et al., 2022) introduces a multilayer annotation model integrating temporal, referential, spatial, and semantic-role information, enabling the annotation of complex narrative structures.

Despite the maturity of temporal annotation schemes, the joint modeling of causal and temporal relations remains relatively rare. These dimensions are strongly interrelated, since causal relations always imply a temporal ordering between events, specifically, that the cause precedes its effect (Tan et al., 2022). However, our investigation confirms that certain cases exhibit alternative temporal configurations, including simultaneity and inclusion, where one event can occur at the same time as the other, encompass, or be encompassed by the other. Some attempts to integrate both causal and temporal dimensions include the annotation

frameworks proposed by Bethard and Martin (2008) and by Mirza and Tonelli (2014). While the former jointly annotates causal and temporal relations between event pairs, the latter combines TimeML-style temporal annotation with lexicalized causal relations based on force-dynamics categories such as *Cause*, *Enable*, and *Prevent*. Nevertheless, most existing approaches still treat causal and temporal information in separate annotation layers.

Some semantic representation formalisms, such as Segmented Discourse Representation Theory (SDRT) (Asher and Lascarides, 2003), Uniform Meaning Representation (UMR) (Gysel et al., 2021), and YARN (Pavlova et al., 2024), adopt a layered architecture in which different semantic phenomena are encoded in distinct but interrelated levels, enabling a structured and modular treatment of meaning. Although these frameworks are not specifically designed to model causal and temporal phenomena in isolation, nor explicitly targeted at causality, their architectures are nonetheless capable of supporting the representation of such relations. In some cases, however, their aim of providing a comprehensive semantic representation leads to considerable complexity, as they incorporate a wide range of linguistic phenomena.

Although temporal annotation frameworks are relatively well established, causal annotation schemes remain less consolidated, more heterogeneous, and in need of broader validation. Moreover, integrated representations that jointly capture causality and temporality are still limited, particularly when considering the need for interoperable frameworks. This gap motivates the development of approaches that model both dimensions simultaneously in order to provide a more comprehensive account of how events are organized in discourse. In this context, our proposal not only integrates these two dimensions, but also adopts an interoperable framework that enables implementation across different semantic formalisms, thus facilitating broader applicability and reuse.

### 3. The process of designing the annotation scheme

The scheme proposed in this work aims to represent causal relations in discourse while simultaneously capturing the temporal relations that hold between the events involved in those relations. The development process involved the integration of two semantic layers, causality and temporality, grounded in the ISO 24617 framework and implemented within the Text2Story semantic annotation architecture. Although the annotation scheme is exemplified with data from European Portuguese, it is, in principle, applicable across languages.

### 3.1. Methodology for developing the annotation scheme

The development of the proposed annotation scheme followed the methodology for the construction and validation of annotation frameworks proposed by [Fernandes et al. \(2025\)](#). This methodology combines conceptual modeling and empirical validation through successive phases, namely: (i) literature review, (ii) design and specification, (iii) empirical validation, and (iv) consolidation and refinement. The initial phase consisted of a comprehensive review of existing approaches to the annotation of causal and temporal relations, including semantic annotation standards and discourse and temporal annotation frameworks. The second phase focused on defining the conceptual structure of the proposed scheme, including its entities, attributes, and relations. The third phase involved pilot annotation experiments on real corpus data, which enabled the systematic identification of inconsistencies, gaps, and incompatibilities. In this phase, the development process followed the MATTER methodology proposed by [Pustejovsky and Stubbs \(2012\)](#). In particular, the iterative subcycle *Model–Annotate* was adopted and operationalized into four stages: *Model*, *Annotate*, *Evaluate*, and *Revise*. This iterative process supported continuous refinement of the scheme based on corpus evidence and annotation feedback. During the development process, this cycle was applied repeatedly because several issues emerged when annotating real data, requiring continuous refinement of the scheme in order to adequately represent causal and temporal relations. As noted by [Silvano et al. \(2021\)](#), building a bootstrapping annotation scheme is a complex and time-consuming task that involves several iterative phases. Finally, a consolidation phase was conducted to refine the model and stabilize the guidelines.

### 3.2. Motivation for using the ISO 24617 framework

The annotation scheme proposed in this work is grounded in the ISO 24617 series of standards for semantic annotation. This framework provides a modular architecture for representing multiple semantic dimensions of discourse, including temporal information (ISO 24617-1: 2012) and discourse relations (ISO 24617-8: 2016). The decision to adopt this framework was motivated by several factors. First, ISO 24617 provides a standardized and interoperable model for semantic annotation, allowing the integration of different semantic layers within a unified structure. Second, the framework has been successfully applied in previous annotation projects, including the Text2Story multilayer annotation scheme proposed by [Silvano et al. \(2021\)](#) and

[Leal et al. \(2022\)](#). Third, the modular architecture of the ISO framework facilitates the integration of discourse-relation and temporal information.

### 3.3. Development of the annotation scheme

The proposed annotation scheme was developed from the Text2Story framework ([Silvano et al., 2021](#)), which has already demonstrated positive results in the annotation of journalistic narratives ([Silvano et al., 2024](#)). Since Text2Story is grounded in the ISO 24617 semantic annotation framework, it provides a suitable basis for extending the representation of narrative structures toward causal information. In this work, the temporal layer inherited from ISO 24617-1: 2012 was preserved with adaptations from Text2Story framework, while a causal layer inspired by ISO 24617-8: 2016 was integrated into the architecture in order to represent both causal and temporal relations within the same annotation scheme.

In our proposal, we adopt a broader notion of *entity* in order to harmonize the temporal and causal layers within a unified annotation model. Following the architectural principles of ISO 24617-1: 2012, the annotation scheme is organized into two complementary structures: *entity structures* and *link structures*. Entity structures contain semantic information associated with a segment of the source text, while link structures represent semantic relations between such segments by establishing links between entity structures. Formally, an annotation structure consists of two sets,  $M$  and  $L$ . The set  $M$  contains pairs  $\langle m, a \rangle$ , where  $m$  corresponds to a markable (a segment of the source text) and  $a$  corresponds to the associated entity structure describing its semantic properties. The set  $L$  contains triples  $\langle r, a_i, a_j \rangle$ , where  $r$  represents a relational annotation and  $a_i$  and  $a_j$  correspond to entity structures participating in that relation. Each element of  $L$  links two entity structures contained in  $M$ , ensuring that relations are defined over annotated textual segments. More complex annotation structures consist of sets of entity structures linked through temporal, causal, and auxiliary relations, thereby allowing the representation of interactions between events, discourse arguments, and explicit causal signals in the text.

Figure 1 summarizes the interaction between the entity structure (events, situations, signals) and the link structure (discourse, temporal, and auxiliary links).

The temporal component of the scheme derives from ISO 24617-1: 2012 (SemAF-Time), which defines a semantic framework for the annotation of time and events. Within this layer, the label *Event* is used to represent eventualities expressed in the

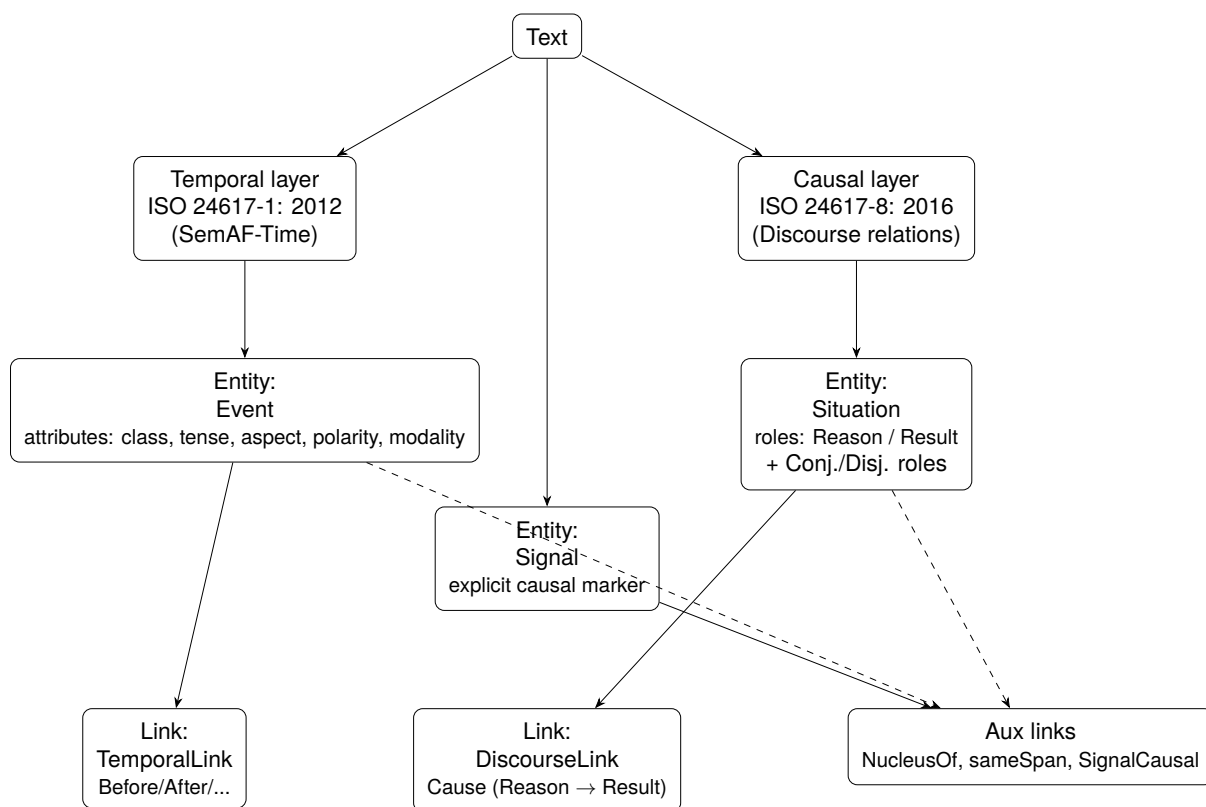


Figure 1: Overall architecture integrating ISO 24617-8: 2016 (causal discourse relations) and ISO 24617-1: 2012 (temporal layer) within Text2Story.

text. Events correspond to occurrences, states, or circumstances and are typically realized by verbs, event nominals, or nominalizations. Each event is characterized through a set of semantic and grammatical attributes, including *Class*, *Type*, *Part of Speech*, *Tense*, following the annotation guidelines of the Text2Story framework (Silvano et al., 2023). Temporal relations between events are represented through temporal links encoding ordering relations such as *Before*, *After*, *Simultaneous*, *Includes*, and *Is\_Included*. These relations allow the chronological organization of the eventualities expressed in the text to be represented.

On the basis of this temporal layer, a causal component was introduced following the discourse-relation framework defined in ISO 24617-8: 2016. This standard provides an interoperable approach to the annotation of discourse relations and establishes correspondences between different theoretical frameworks, including Rhetorical Structure Theory (RST) (Mann and Thompson, 1988), Segmented Discourse Representation Theory (SDRT) (Asher and Lascarides, 2003), the Penn Discourse Treebank (PDTB) (Prasad et al., 2008), Hobbs’ Theory of Discourse Coherence (HTDC) (Hobbs, 1985), and the Cognitive Approach to Coherence Relations (CCR) (Sanders et al., 1992). In the ISO framework, discourse relations are defined as semantic links between two situations, which

may correspond to clauses, sentences, nominalizations, or larger discourse segments. To incorporate causal relations into the proposed scheme, a new entity label called *Situation* was introduced in the entity structure. This label allows the annotation of situations expressed in texts as clauses, sentences, noun phrases, nouns, or verbs. Within *Situation*, the two core semantic roles proposed by ISO 24617-8: 2016 were adopted: *Reason* and *Result*. These roles correspond to the two arguments of the causal relation and enable the explicit identification of the textual segments that function as cause and effect. To establish the causal relation between these arguments, a new link type called *Discourse* was introduced in the link structure. Thus, according to ISO 24617-8:2016, the relevant discourse relation is labeled *Cause* and involves two arguments associated with distinct semantic roles, namely *Reason* and *Result*<sup>1</sup>, where these two terms are interpreted as having the same underlying meaning but different functions: one as a discourse relation and the other as a semantic role within that relation.

The scheme also includes the annotation of ex-

<sup>1</sup>Although one may argue that a distinction may be drawn between ‘cause’ and ‘reason’ in terms of their semantic contribution, this work adopts the terminology and conceptualization provided by ISO 24617-8:2016.

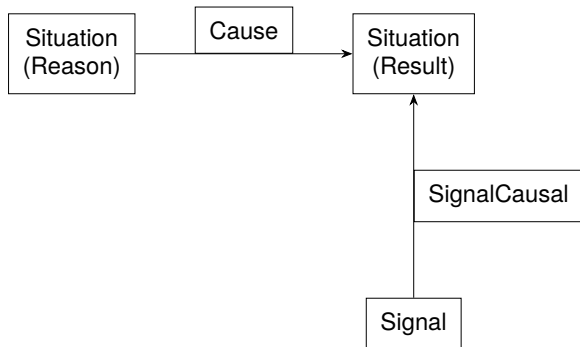


Figure 2: Causal layer: *Cause* links reason and result situations; explicit *Signal* is associated with the result argument.

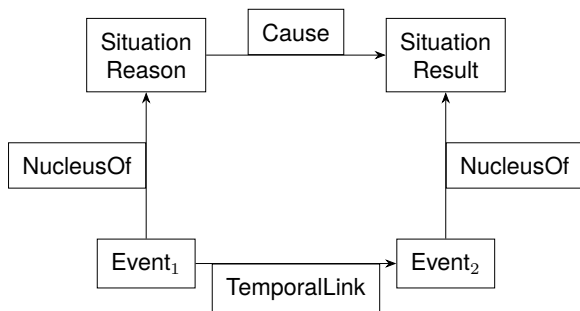


Figure 3: Integration: causal relations link *Situations*, while temporal relations are encoded between the main *Events* associated with each argument.

explicit causal markers. For this purpose, a label called *Signal* was introduced to identify lexical or multiword expressions that explicitly indicate a causal relation between two arguments. Based on previous work, particularly the Penn Discourse Treebank and linguistic studies on causal markers, explicit signals were categorized into grammatical classes. A specific link type, also called *Signal*, was added to the link structure to associate the signal with the result argument of the causal relation.

Figure 2 summarizes the core representation of the causal layer, including the association of an explicit signal with the result argument via a dedicated link type.

Figure 3 illustrates the adopted integration strategy.

The integration of the temporal and causal layers required several modeling decisions. In the ISO framework, causal relations are defined between arguments, whereas temporal relations are defined between events. However, the standard does not explicitly define how these two layers should interact when both causal and temporal information must be represented simultaneously. This limitation becomes particularly evident when applying the annotation scheme to real data, as illustrated

in the following example<sup>2</sup>.

- (1) *O fogo provocou danos no recheio e na estrutura do edifício.* (Lusa2\_36)  
*The fire caused damage to the interior and the structure of the building.*

In this example, the sentence clearly expresses both a causal relation and a temporal ordering between the events involved. To address this issue, the proposed scheme establishes temporal relations between the main events associated with each causal argument rather than directly between the arguments themselves. Thus, the causal relation is maintained between the *Reason* and the *Result*, while the temporal relation is encoded between the events that instantiate these arguments. To support this interaction between layers, a relation called *NucleusOf* was introduced. This relation links an *Event* to the *Situation* to which it belongs, making it possible to explicitly identify the main event associated with each causal argument.

During the annotation experiments conducted on a set of news articles from the Lusa news agency, several issues emerged that required adjustments to the scheme. One major challenge concerned the representation of complex causal configurations. Although ISO 24617-8: 2016 defines causal relations as binary relations between a reason and a result, the corpus revealed more complex structures. In some cases, a single reason leads to multiple results, whereas in other cases several reasons converge toward a single result.

The following examples illustrate these situations:

- (2) *Um fenómeno meteorológico geralmente associado ao aumento das temperaturas, a secas em algumas partes do mundo e a chuvas fortes noutras.* (Lusa2\_100)  
*A meteorological phenomenon generally associated with rising temperatures, droughts in some parts of the world, and heavy rains in others.*
- (3) *Pelo menos quatro pessoas continuam hoje desaparecidas e cerca de 20 mil foram afetadas por inundações no sul da Tailândia, na sequência de uma forte tempestade.* (Lusa2\_116)  
*At least four people remain missing today and around twenty thousand have been af-*

<sup>2</sup>For illustrative purposes, we use European Portuguese data collected from the Lusa news agency, the largest Portuguese-language news agency. The examples presented in this paper are drawn from a dataset currently under preparation for publication. The original Portuguese sentences are provided alongside their English translations, and each example is identified by its corresponding reference (e.g., Lusa2\_36).

ected by floods in southern Thailand following a strong storm.

- (4) *As outras mortes foram causadas pela queda de ramos ou desmoronamento de estruturas.* (Lusa2\_109)  
*The other deaths were caused by falling branches or collapsing structures.*

To represent these configurations while preserving the ISO requirement that causal relations involve only two arguments, four additional roles were introduced within the *Situation* entity: *Conjunction1*, *Conjunction2*, *Disjunction1*, and *Disjunction2*. These roles instantiate the ISO definitions of the Conjunction relation (both arguments hold: Arg1 Arg2) and the Disjunction relation (non-exclusive alternatives: Arg1 Arg2). Specifically, *Conjunction1* and *Conjunction2* capture multiple elements converging on a shared situation (multiple reasons → one result; one reason → multiple results), while *Disjunction1* and *Disjunction2* capture alternative elements (multiple alternative reason → one result; one reason → multiple alternative results). This approach allows coordinated or alternative causal elements to be represented as components of a single argument, while still allowing more complex causal configurations to be represented.

A further difficulty concerns the annotation of discontinuous spans due to limitations of the INCEPTION annotation tool, used in this project, which does not allow non-contiguous segments to be annotated as a single markable. This situation occurs in examples such as:

- (5) *a morte por electrocussão de um homem de 50 anos* (Lusa2\_109)  
*the death by electrocution of a 50-year-old man*

In this example, the argument with the role of *Result* includes both *a morte* and *de um homem de 50 anos*. To preserve the semantic unity of the argument, a relation called *sameSpan* was introduced to link the discontinuous segments belonging to the same argument.

Finally, difficulties also arose in identifying the events associated with each causal relation in sentences containing multiple events. For instance:

- (6) *Cerca de 1.000 pessoas ficaram feridas e mais de 14 mil casas foram destruídas na sequência do sismo mais mortífero dos últimos nove anos na China.*  
*Around 1,000 people were injured and more than 14,000 houses were destroyed following the deadliest earthquake in China in the last nine years.*

In this case, the events *feridas* and *destruídas* must be explicitly associated with the correspond-

ing causal argument. The introduction of the *NucleusOf* relation made it possible to link these events to their corresponding *Situation* argument and thus to determine their role in the causal relation.

#### 4. CATS: Causal and Temporal Annotation Scheme

After several rounds of experimentation and refinement, the final version of the annotation scheme was established. Figure 4 presents the complete CATS scheme: an annotation scheme of causality and temporal structure.

As shown in Figure 4, the entity structure includes the labels *Event*, *Situation*, and *Signal*, each associated with a set of attributes and values. The link structure includes the relations *SameSpan*, *NucleusOf*, two technical links, and *Signal*, and *Discourse*, the latter encompassing both causal and temporal relations.

The newly proposed label *Situation* tag corresponds to the arguments of discourse relations and is realized by simple or complex linguistic expressions such as clauses, nominalizations, phrases, or sets of phrases. The roles associated with *Situation* are the following: *Reason* and *Result*, corresponding to the two arguments of a causal relation; *Conjunction1* and *Conjunction2*, used to represent coordinated causes or effects; *Disjunction1* and *Disjunction2*, used to represent alternative causal configurations.

The *Signal* label refers to lexical items or multiword expressions that explicitly mark a causal relation. In the annotation scheme, signals are not considered part of the causal arguments but function exclusively as markers linking the arguments. The attributes associated with this label correspond to different grammatical categories of causal signals: *Verb* - verbs that encode causality in their semantics (e.g., *provoke*, *cause*, *lead to*); *Noun* - nouns referring to the cause or origin of an event (e.g., *cause*, *origin*, *motive*); *Conjunction* - conjunctions and conjunctive expressions linking clauses to express cause or explanation (e.g., *because*, *since*, *as*); *Preposition* - prepositions or prepositional expressions introducing causal relations (e.g., *due to*, *because of*, *thanks to*); *Adverb* - adverbs or adverbial expressions expressing causal relations (e.g., *thus*, *consequently*).

The *Event* tag represents occurrences, states, or circumstances described in the text. Following ISO 24617-1: 2012 and the annotation manual by Silvano et al. (2023), events are characterized by several attributes: *Class* - occurrence, state, reporting, perception, aspectual, IState, IAction; *Type* - aspectual type (state, process, or transition); *Pos* - grammatical category (verb, noun, adjective, or preposition); *Tense* - temporal location of the event

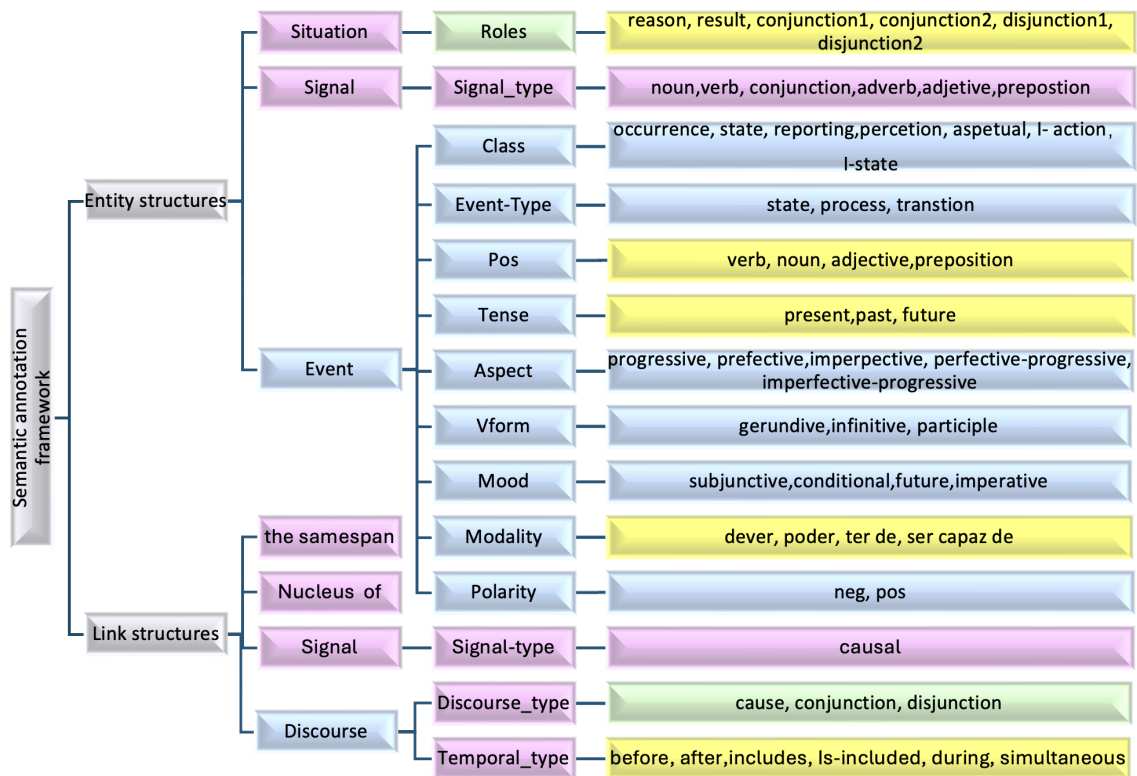


Figure 4: CATS: Annotation scheme for causality and temporal structure. Color coding: blue indicates labels/attributes from ISO 24617-1:2012; green indicates labels/attributes from ISO 24617-8:2016; yellow indicates labels/attributes adapted from either standard; pink indicates newly proposed labels/attributes.

(present, past, future); *Aspect* - aspectual value (progressive, perfective, imperfective, imperfective progressive); *VForm* - non-finite verb forms (gerund, infinitive, participle); *Mood* - verbal mood (subjunctive, conditional, future); *Modality* - modal value (e.g., *dever* 'must', *poder* 'can'); *Polarity* - positive or negative.

The link structure includes relations representing both semantic and technical connections between annotated elements. The *Discourse* relation includes both causal and temporal relations. A causal relation is established when one argument provides an explanation for another. In the annotation scheme, the causal relation connects a *Reason* argument to a *Result* argument, with directionality from *Reason* to *Result*. Temporal relations represent the ordering between events associated with the causal arguments. The values of temporal relations include: *Before* - one event precedes another; *After* - one event follows another; *Simultaneous* - two events occur at the same time; *Includes* - one event temporally includes another; *IsIncluded* - one event is temporally included in another.

In addition to these semantic relations, the scheme includes three technical relations: *Signal* - links a causal signal to the corresponding causal relation, connecting the signal to the *Result* argu-

ment; *NucleusOf* - links events to the situations (*Reason* or *Result* arguments) in which they occur; *the sameSpan* - connects non-contiguous segments belonging to the same causal argument.

In our proposal, the annotation procedure should be the following:

1. Identification of textual segments instantiating a causal relation;
2. Annotation of the *Reason* and *Result* arguments;
3. Identification and annotation of the main events of the *Reason* and *Result* arguments;
4. Identification and annotation of *causal signals*;
5. Annotation of *NucleusOf* relations linking events to arguments;
6. Annotation of *causal* and *temporal relations*;
7. Annotation of *signal relations*.

To identify causal relations, we draw on five tests proposed by (Grivaz, 2010; Dunietz et al., 2017): *Why* test, *Temporal Order* test, *Counterfactual* test, *Ontological Asymmetry* test, and *Linguistic* test. The *Why* test assesses whether the effect can be

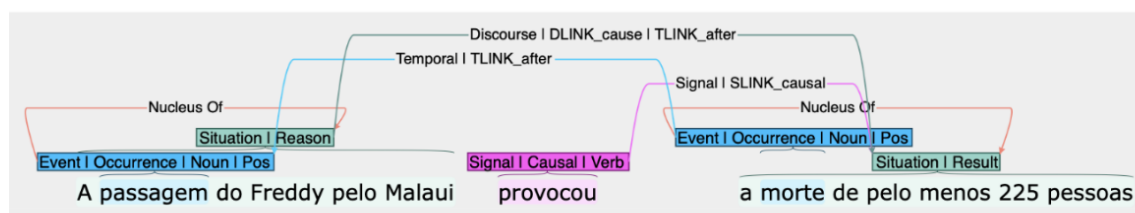


Figure 5: Annotated example according to the CATS scheme.

explained by asking why it occurred; *Temporal Order* checks whether the cause precedes the effect; *Counterfactual* test evaluates whether the effect would not occur in the absence of the cause; *Ontological Asymmetry* captures the non-symmetrical nature of causality; and *Linguistic* verifies whether the relation can be paraphrased as "It is because (of) X that Y" or "X causes Y." In this work, we adopt the *Why*, *Ontological Asymmetry*, and *Linguistic* tests as the most relevant. The *Temporal Order* test is excluded, as causal relations may involve simultaneity or posteriority. The *Counterfactual* test is also excluded, since the notion of causality adopted in ISO 24617-8:2016 is broader and does not always entail strict counterfactual dependence.

Example (7) illustrates the annotation of a causal relation in a news sentence.

- (7) *A passagem do Freddy pelo Malawi provocou a morte de pelo menos 225 pessoas. Freddy's passage through Malawi caused the death of at least 225 people.*

The annotation of this sentence includes the following elements:

- **Reason argument:** "A passagem do Freddy pelo Malawi" (Freddy's passage through Malawi);
- **Result argument:** "a morte de pelo menos 225 pessoas" (the death of at least 225 people);
- **Events:** "passagem" (passage) and "morte" (death);
- **Causal signal:** "provocou" (caused);
- **NucleusOf:** linking the event "passagem" to the *Reason* argument and the event "morte" to the *Result* argument;
- **Causal and temporal relations:** the *Reason* argument is linked to the *Result* argument through the relation *Cause*, and a temporal relation *After* is established between the corresponding events;

- **Signal relation:** linking the signal "provocou" to the *Result* argument.

Figure 5 presents an example of annotation carried out in the INCEpTION tool (Klie et al., 2018) using the CATS annotation scheme.

## 5. Conclusion

This article presents a new annotation framework designed to jointly capture causal and temporal information. To achieve this objective, we integrate elements from the temporal layer defined in ISO 24617-1: 2012 and the discourse relation framework of ISO 24617-8: 2016, resulting in CATS, an annotation scheme of causality and temporal structure.

Building on the Text2Story annotation framework, the proposed model extends the entity structure with additional labels and attributes for representing events, situations, and causal signals, while the link structure incorporates relations required to encode causal and temporal dependencies between annotated elements. This integration enables a coherent representation of the interaction between causal relations and the temporal ordering of events.

CATS therefore provides a unified framework for the integrated annotation of causality and temporality, while remaining compatible with the broader ISO 24617 architecture. In addition to causal relations, the framework also supports the annotation of other discourse relations defined in ISO 24617-8: 2016.

Future work will focus on presenting the results of applying the proposed annotation scheme to a dataset of Portuguese news articles to provide evidence of its effectiveness in real annotation scenarios. Furthermore, we plan to extend the scheme by incorporating additional discourse relations, thereby broadening its applicability for the semantic annotation of complex discourse structures.

## 6. Acknowledgements

This work is 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 10.54499/2022.09312.PTDC).

## 7. Bibliographical References

- Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.
- Steven Bethard, William Corvey, Sara Klingenstein, and James H. Martin. 2008. Building a corpus of temporal-causal structure. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC 2008)*, pages 908–915. LREC 2008.
- Steven Bethard and James H. Martin. 2008. Learning semantic links from a corpus of parallel temporal and causal relations. In *Proceedings of ACL-08: HLT*, pages 177–185. Association for Computational Linguistics.
- Harry Bunt and Rashmi Prasad. 2016. *ISO DR-core (ISO 24617-8): Core concepts for the annotation of discourse relations*. In *Proceedings of the Joint ACL-ISO Workshop on Interoperable Semantic Annotation*.
- Francisco Costa and António Branco. 2012a. *Extracting temporal information from Portuguese texts*. In Helena Caseli, Aline Villavicencio, António Teixeira, and Fernando Perdigão, editors, *Computational Processing of the Portuguese Language*, volume 7243 of *Lecture Notes in Computer Science*, pages 105–116. Springer, Berlin, Heidelberg.
- Francisco Costa and António Branco. 2012b. TimeBankPT: A TimeML annotated corpus of Portuguese. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*, pages 3727–3734, Istanbul, Turkey. European Language Resources Association (ELRA).
- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2015. *Annotating causal language using corpus lexicography of constructions*. In *Proceedings of the 9th Linguistic Annotation Workshop*, pages 188–196.
- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2017. *The BECauSE corpus 2.0: Annotating causality and overlapping relations*. In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 95–104.
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2022. *Causal inference in natural language processing: Estimation, prediction, interpretation and beyond*. *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- 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.
- Roxana Girju, Preslav Nakov, Vivi Nastase, Stan Szpakowicz, Peter Turney, and Deniz Yuret. 2007. *SemEval-2007 task 04: Classification of semantic relations between nominals*. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 13–18.
- Cécile Grivaz. 2010. *Human judgements on causation in French texts*. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Jens E. L. Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah R. Moeller, Jiarui Yao, Tim O’Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, Jan Hajic, James H. Martin, Stephan Oepen, Martha Palmer, James Pustejovsky, Rosa Vallejos, and Nianwen Xue. 2021. *Designing a uniform meaning representation for natural language processing*. *Künstliche Intell.*, 35(3):343–360.
- Jerry R. Hobbs. 1985. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information (CSLI), Stanford University.
- ISO. 2012. ISO 24617-1: 2012, Language Resource Management - Semantic Annotation Framework (SemAF) - Part 1: Time and events. Technical report, International Organisation for Standardisation ISO, Geneva.
- ISO. 2016. ISO 24617-8: 2016, Language Resource Management - Semantic Annotation

- Framework (SemAF) - Part 8: Semantic RRelations in Discourse (DR-Core). Technical report, International Organisation for Standardisation ISO, Geneva.
- Andrew Kehler. 2002. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications.
- Jan-Christoph Klie, Michael Bugert, Beto Boulosa, 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, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Mahnaz Koupaee, Xueying Bai, Mudan Chen, Greg Durrett, Nathanael Chambers, and Niranjana Balasubramanian. 2025. Causal graph based event reasoning using semantic relation experts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26169–26199.
- Alex Lascarides and Nicholas Asher. 1993. [Temporal interpretation, discourse relations and commonsense entailment](#). *Linguistics and Philosophy*, 16(5):437–493.
- 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.
- Inderjeet Mani and Barry Schiffman. 2004. Temporally anchoring and ordering events in news.
- William C. Mann and Sandra A. Thompson. 1988. [Rhetorical structure theory: Toward a functional theory of text organization](#). *Text*, 8(3):243–281.
- Paramita Mirza and Sara Tonelli. 2014. [An analysis of causality between events and its relation to temporal information](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2097–2106, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Marc Moens and Mark Steedman. 1988. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28.
- Qiang Ning, Zhili Feng, Hao Wu, and Dan Roth. 2018a. [Joint reasoning for temporal and causal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2278–2288.
- Qiang Ning, Hao Wu, and Dan Roth. 2018b. [A multi-axis annotation scheme for event temporal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1318–1328, Melbourne, Australia. Association for Computational Linguistics.
- Siyana Pavlova, Maxime Amblard, and Bruno Guillaume. 2024. [YARN is all you knit: Encoding multiple semantic phenomena with layers](#). In *Proceedings of the Fifth International Workshop on Designing Meaning Representations @ LREC-COLING 2024*, pages 66–76, Torino, Italia. ELRA and ICCL.
- Judea Pearl. 2009. *Causality: Models, Reasoning, and Inference*, 2 edition. Cambridge University Press, Cambridge.
- Rashmi Prasad and Harry Bunt. 2015. [Semantic relations in discourse: The current state of ISO 24617-8](#). In *Proceedings of the 11th Joint ACL-ISO Workshop on Interoperable Semantic Annotation (ISA-11)*.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltasakaki, Livio Robaldo, Aravind K. Joshi, and Bonnie Webber. 2008. The Penn Discourse Treebank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC 2008)*, pages 2961–2968, Marrakech, Morocco. European Language Resources Association.
- James Pustejovsky, Jose M. Castano, Robert Inghria, Roser Sauri, Robert Gaizauskas, Andrea Setzer, and Graham Katz. 2003a. [TimeML: Robust specification of event and temporal expressions in text](#). In *Proceedings of the 5th International Workshop on Computational Semantics (IWCS-5)*.
- James Pustejovsky, Patrick Hanks, Roser Sauri, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, and Marcia Lazo. 2003b. The timebank corpus. In *Proceedings of Corpus Linguistics 2003*, pages 647–656.
- James Pustejovsky, Kiyong Lee, Harry Bunt, and Laurent Romary. 2010. [ISO-TimeML: An international standard for semantic annotation](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*

- (LREC'10), Valletta, Malta. European Language Resources Association (ELRA).
- James Pustejovsky and Amber Stubbs. 2012. *Natural Language Annotation for Machine Learning - a Guide to Corpus-Building for Applications*. O'Reilly.
- Ted Sanders, Wilbert Spooren, and Leo Noordman. 1992. [Toward a taxonomy of coherence relations](#). *Discourse Processes*, 15:1–35.
- Andrea Setzer. 2001. *Temporal Information in Newswire Articles: an Annotation Scheme and Corpus Study*. Ph.D. thesis, University of Sheffield.
- 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, Alípio Mário Jorge, António Leal, Evelin Amorim, Hugo Sousa, Inês Cantante, Ricardo Campos, and Sérgio Nunes. 2023. [Text2story lusa annotated corpus](#). Dataset.
- 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.
- Steven A. Sloman. 2005. *Causal Models: How People Think about the World and Its Alternatives*. Oxford University Press.
- Leonard Talmy. 2000. *Toward a Cognitive Semantics, Volume II: Typology and Process in Concept Structuring*. MIT Press.
- Fiona Anting Tan, Liyuan Lee, and Yejin Choi. 2022. Causal news corpus: Annotating causality in event sentences from news. In *Proceedings of the 29th International Conference on Computational Linguistics (COLING)*, pages 2298–2308, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- James Woodward. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
- Jie Yang, Soyeon Caren Han, and Josiah Poon. 2022. [A survey on extraction of causal relations from natural language text](#). *Knowledge and Information Systems*, 64:1161–1186.