

Universal Discourse Relations: A Proposal

Anna Latusek Maciej Ogrodniczuk Alina Wróblewska Bartosz Żuk
Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland
{a.latusek, m.ogrodniczuk, alina, b.zuk}@ipipan.waw.pl

Abstract

This paper introduces a novel ‘universal’ approach to discourse annotation, serving as a comprehensive synthesis of the ISO 24617-8 semantic annotation framework and a newly developed multilayer model of coherence relations. To address the complexities of text analysis, we present a hierarchical classification and a systematic decision tree. By unifying disparate formalisms, our model provides researchers with a robust, standardised methodology for analysing complex discourse structures across various linguistic contexts.

1 Introduction

Discourse relations, also known as coherence or rhetorical relations, are semantic links between text segments that provide structure and coherence to a discourse. Even though there exists a multitude of formalisms for discourse relation representation, none of them builds on differences in their nature, most often treating coherence-building, temporal, referential, meta-textual, structure-oriented and function-oriented relations as notions from the same repertoire, merging them in a common schema. Aiming to capture the multifaceted essence of text structure with the document-level coherence in mind, we separate these various phenomena into independent layers with discourse relations as the backbone.

We drew on the ISO framework (ISO, 2016). Its aim was to confront different discourse-annotation formalisms and extract what recurs across all of them: a universal set of relations that appears both in shallow approaches (such as PDTB) and in deep approaches (such as RST). By adopting this umbrella perspective – moving from the specific to the general – it becomes possible to attempt a unified multilingual description. And this is precisely our main objective: to describe discourse cross-linguistically by applying general annotation

principles that can be used across different language types, rather than being tailored to a single one.

In this paper, we propose: 1) a multilayer representation of discourse, taking into account these various aspects of its organisation, structure, referential properties, ways of attribution etc., 2) its discourse-centered practical approach offering a flat representation of the coherence relations, most similar to what many formalisms previously attempted and 3) the decision tree used to annotate universal relations according to the proposed model.

2 Related Work

2.1 Multilevel Modelling of Discourse Analysis

Discourse can be analyzed from structural/formal perspective, including frameworks such as RST (Mann and Thompson, 1987) and SDRT (Asher and Vieu, 2005), as well as from a lexically grounded, relation-based perspective, as in PDTB (Prasad et al., 2008), with additional influence from functional (Halliday and Hasan, 1976; Halliday, 2014), pragmatic (Grice, 1975), and cognitive (van Dijk, 1980; Kintsch, 1998) approaches.

All of these theories share a common principle: in order to consider a string of sentences as discourse, its coherence is crucial. Achieving coherence requires attention at multiple levels, often extending beyond discourse structure to the level of the document as a whole. The multifaceted and multilevel nature of discourse description has been addressed, among others, in the multilingual context by Webber et al. (2012) and Jurafsky and Martin (2026), and in the Polish context by Żydek-Bednarczuk (2005). Multilevel discourse annotation is also applied in practice in some corpora. One of the well-known examples is the Potsdam Commentary Corpus (Stede and Neumann, 2014),

which provides space for annotating both PDTB-style discourse relations and text structure drawn from RST. The multilayered nature of discourse is also discussed using as exemplification specific types of texts, such as dialogues (Petukhova et al., 2011).

Studying and describing discourse relations – also referred to in the literature as rhetorical (e.g., in RST or SDRT) or coherence relations (Hobbs, 1979) – therefore requires a broader perspective that views discourse as a multilevel structure.

On the other hand, building a multilayered model of discourse requires a detailed analysis of each layer separately. Since our first goal was to develop annotation guidelines for coherence relations and their composition into coherent thematic segments, we focus on how existing formalisms – and attempts to unify them – address the various aspects of coherence-building relations and the identification of topic structure.

2.2 Coherence Relations

A discourse relation is considered to be a connection between two text segments (so-called Discourse Units – DUs) that contributes meaning to the utterance: it indicates dependencies between the segments, extends the narrative (or, following ISO, the situation), and situates it within a context.

ISO 24617-8 seeks to provide a pivot discourse relations representation by defining a set of 20 core relations that are commonly found across existing frameworks. Two main types of relations are distinguished: symmetric and asymmetric, and they hold between segments referred to as arguments A and B. Symmetry means that (REL, A, B) is equivalent to (REL, B, A).

Unification of relations across different approaches is also proposed in DISRPT (Braud et al., 2025). The main motivation is to harmonise datasets annotated according to different frameworks and collected across successive editions of the shared tasks. DISRPT largely corresponds to ISO, with a core set of 17 relations.

In RST rhetorical relations are organised into a hierarchical tree structure, where segments are connected through nucleus–satellite (asymmetric) or multinuclear (symmetric) configurations, reflecting their relative importance. In total, they distinguish 78 relations, enclosed in 16 classes. Two main types of relations are subject-matter and presentational relations, which concentrate on what the speaker/writer wants to express. The most signifi-

cant reflection of this theory is the RST Treebank (Carlson et al., 2001).

The foundation of SDRT (Lascarides and Asher, 2008) is also formed by rhetorical relations, which can be either subordinating (hierarchical) or coordinating (equal in status), and whose descriptions are grounded in formal logic. Since multiple relations can exist between discourse segments, these dependencies are represented as graphs.

PDTB adopts a lexically grounded approach centred on discourse connectives, which means that relations are annotated locally between pairs of arguments, typically anchored in explicit or implicit discourse connectives, without constructing a full global structure. Relations are grouped into four major classes: temporal, contingency, comparison and expansion, in which both semantic and pragmatic sources of coherence can be found. The most up-to-date, standardised version, as of today, is PDTB 3.0 (Prasad et al., 2014).

Pragmatic Relations Although ISO 24617-8 places considerable emphasis on the semantic dimension of discourse relations, most contemporary formalisms acknowledge that pragmatic factors are inseparable from how coherence is constructed. ISO 24617-5 (ISO, 2014) explicitly defines a discourse relation as a “semantic/pragmatic relation that holds among two or more circumstances” (p. 2), underscoring that interpretation goes beyond the literal content of the segments involved. This pragmatic layer is also visible in established frameworks: in RST, presentational relations capture the speaker’s communicative intentions, while in the PDTB taxonomy several relations are explicitly categorised as pragmatic (e.g., *Pragmatic Cause*, *Pragmatic Contrast*, which encode the writer’s intended argumentative or rhetorical effect rather than objective semantic relations).

In practice, discourse interpretation cannot be cleanly separated from pragmatic inference, a point repeatedly emphasised in the literature (Hobbs, 1979; Asher and Lascarides, 2003). Speaker intention, relevance, and communicative goals play a central role in shaping coherence. This is particularly evident in dialogue, where pragmatic dependencies are often modelled through speech act-based approaches (Petukhova et al., 2011), as formalised in ISO 24617-2 (ISO, 2020).

A separate, yet equally significant line of inquiry concerns the pragmaticalization of connectives (Degand and Evers-Vermeul, 2015; Zufferey

and Degand, 2024), that is, discourse relation markers which form the basis of surface-oriented approaches to discourse.

Temporal Relations ISO 24617-8 builds its set of relations on the basis of classical theories of discourse, incorporating temporal relations into the inventory of discourse relations. At the same time, it allows for the annotation of the same segments with more than one relation. It is closely linked to dedicated temporal annotation frameworks such as TimeML, which model time as a separate layer spanning the entire discourse (Pustejovsky et al., 2003). Another argument supporting the view that temporal relations constitute a separate layer of discourse coherence markers can be found in the work of the developers of SDRT (Asher and Lascarides, 1993), where temporal relations are treated as a derivative of discourse relations and commonsense reasoning.

Attribution *Attribution* relation, linking the attribution phrase and the attributed content, is treated differently across discourse frameworks: in PDTB it forms a separate annotation layer linked to the arguments of a relation, whereas in RST, SDRT, GraphBank (Wolf et al., 2005) it appears as a dedicated rhetorical relation structuring the text. However, ISO 24617-8 states that this is “perhaps motivated by the need to not leave any part of the text unconnected” (ISO, 2016, p. 8), and it also excludes attribution in its core set of discourse relations, recommending it instead as a separate layer of annotation. Potter (2019) does the same, arguing that attribution should not be treated as a standard rhetorical relation but rather as a separate discourse annotation layer that introduces a distinct level of perspectival and source-based interpretation. This specific nature of attribution relations is best captured in the PARC corpus, which was created precisely for the purpose of annotating attribution (Pareti, 2016).

Referential Relations Existing discourse representation formalisms take one of the two approaches to reference (entity-based) relations: either by including them in the set of discourse relations, based on the assumption that one discourse unit provides additional or background information about the entity from another unit, or by separating them from discourse relations completely.

RST does not provide a formal mechanism for tracking entity-level anaphora; instead, the frame-

work prioritises characterising text structure primarily in terms of functional and hierarchical relations that hold between text parts.

PDTB maintains a strict separation between entity coreference and discourse relations, labelling relation between corresponding arguments as Ent-Rel (entity-based coherence) when no stronger discourse relation can be inferred¹.

ISO 24617-8 specification “provides for only a single relation to capture these entity-based relations, called *Expansion*. In the future, the second part of the project, of which DR-core forms the first part, should clarify this relation and capture more fine-grained distinctions” (ISO, 2016, p. 10).

2.3 Topic Structure

Topic structure is a central aspect of discourse organisation and has been studied from multiple perspectives in computational linguistics. It has been particularly researched within theories of discourse segmentations. Hearst (1997) introduced Text Tiling, an algorithm for segmenting texts into multi-paragraph units based on lexical cohesion. These units often correspond to topic shifts. Similarly, Beeferman et al. (1999) proposed statistical approaches to linear text segmentation, showing that topic boundaries can be identified using distributional properties of language. These approaches are effective for identifying topic segment boundaries using surface-level features. However, they do not model internal organisation and coherence of topic segments.

Furthermore, topic modelling techniques, such as Latent Dirichlet Allocation (LDA, Blei et al., 2003), offer a complementary perspective by identifying latent thematic structure in texts, i.e. detecting topic shifts and grouping text segments into coherent thematic units. However, this approach represents topics as probabilistic distributions over words and does not explicitly capture their role in structuring discourse units or their interaction with other layers of discourse representation. Therefore, topic structure should be represented explicitly as a separate discourse layer that not only groups discourse units into thematically coherent segments but also reflects their role in structuring the text.

¹In PDTB “EntRel, while being a semantic relation, is not considered a discourse relation with situations as arguments, and is therefore not included in the classification” (ISO, 2016, p. 27).

3 A Multilevel Model of Discourse Analysis

We introduce a novel multilevel model of discourse analysis. The proposed model is designed to capture the complexity of discourse by integrating multiple layers of representations within a unified framework. It enables the systematic analysis of relations both within and across discourse levels and offers a comprehensive account of text coherence and organisation.

We grounded our work in the ISO 24617-8 standard (ISO, 2016), using the set of coherence relations established there, while proposing certain adjustments, shifts, and extensions to that core inventory. Furthermore, we assume that a text cannot be viewed through the lens of only a single layer of coherence (*coherence layer*²). The complex structure of discourse comprises multiple layers, which not only requires a multifaceted description but, above all, precise differentiation among these layers during the annotation process. Viewing discourse as a big entity anchored in a larger whole, we add additional layers to our metamodel: *metatext layer*, *topic layer*, *temporal layer*, and *referential layer*. That provides a complete picture of the document relations that, at different levels, contribute to its coherence. The structure and key components of the model are presented in Figure 1 and outlined below.

Coherence relations Within the set of coherence relations we adopted, four sublevels are distinguished: *discourse*, *attribution*, *pragmatic*, and *formatting* relations. The discourse and pragmatic levels are further fine-grained into more detailed categories (see Section 4 for a detailed description of these types).

It is necessary to distinguish between intra- and inter-sentential coherence. Intra-sentential coherence concerns structural and grammatical relations, as well as internal consistency of discourse units, such as sentences, clauses, and standalone (verb) phrases (Ogrodniczuk et al., 2025). A discourse unit is considered intra-sententially coherent when it hangs together, i.e., all of its components interact to produce a clear and interpretable meaning. Intra-sentential coherence depends on several factors, including grammatical well-formedness (e.g., *I am*

²The term *coherent* is used intentionally as the most appropriate alternative to *semantic*. At the same time, we acknowledge that relations at other layers also contribute to the overall coherence of the text.

reading vs. **I is reading*; *He went home* vs. **He went chess*), temporal consistency (e.g., *I didn't sleep yesterday* vs. **I don't sleep yesterday*), and referential consistency (e.g., *Maria wstała i wyszła* [Maria got up_{fem} and left_{fem}] vs. **Maria wstała i wyszedł* [Maria got up_{fem} and left_{masc}]). These intra-sentential relations are not represented in our model schema in Figure 1.

Inter-sentential coherence, in turn, considers semantic relations between discourse units. These relations are non-structural in nature and contribute to the overall organisation and interpretability of the text. They capture how segments can be linked and how the interpretation of each segment depends on its relation to others. When a sequence of segments is connected through such semantic relations, it forms a coherent text, i.e., a discourse.

For illustration, we present an excerpt from a Sherlock Holmes story (Doyle, 1927), annotated with selected relations from each layer (see Figure 2). The *discourse* layer is illustrated by the relation *Entity-Elaboration*. In addition, the excerpt includes an instance of the *Attribution* relation, the *Question-Answer* relation from the *pragmatic* subset, and three examples of the *Formatting* relations. These simplified annotations demonstrate how different layers contribute to the overall structure and interpretation of the text.

Topic layer The topic layer represents a higher level of discourse organisation and consists of text passages composed of discourse units connected by coherence relations. Each passage is unified by a common topic, forming a thematically coherent segment of the text. Individual topic segments are thus locally (topically) coherent (Jurafsky and Martin, 2026) and are lexically cohesive (Halliday and Hasan, 1976). The topic layer, thus, captures how individual, consecutive discourse units are grouped into larger structures that reflect topical continuity and shifts within the discourse. Topic segments are genre-dependent; they may correspond to entire sections (e.g., Related Work, Results) in scientific articles, to scenes or descriptive passages in narratives, or to entire news articles, etc. Topic segments should cover the entire discourse, except for meta-textual passages, and they constitute a cohesive plot, document or article structure (e.g., a scientific article may not include two independent Related Work sections).

In Figure 2, six topic segments are identified and marked with red frames: (1) reading the first letter,

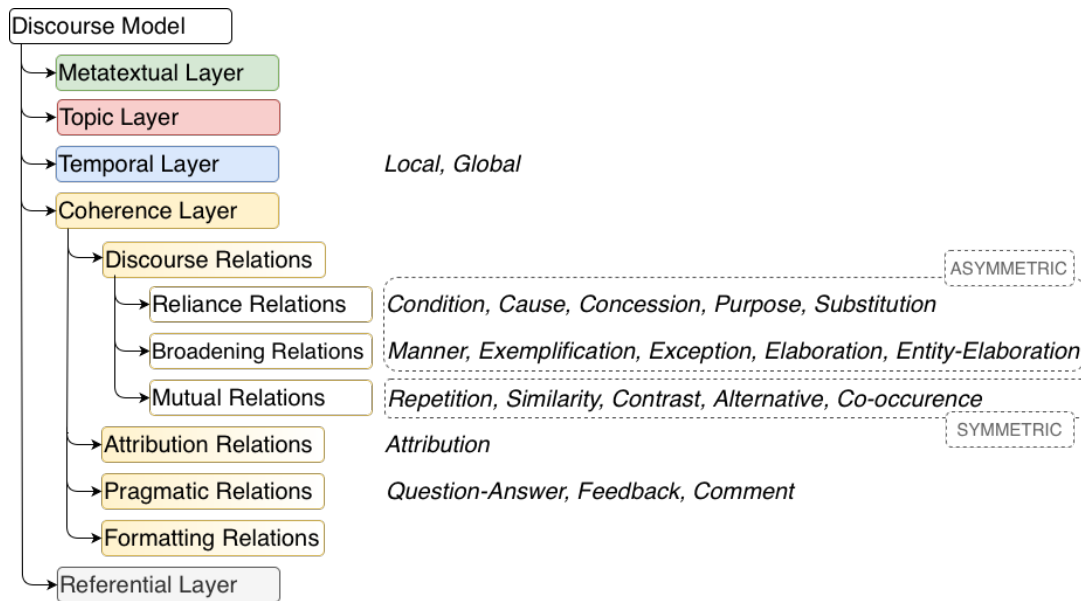


Figure 1: Multilayer discourse model

(2) the first letter, (3) discussing the letter, (4) reading the index volume, (5) discussion of vampires, and (6) discussion based on the second letter.

Temporal layer The temporal layer organises discourse with respect to time and complements the other layers by explicitly representing temporal information. We model it along two dimensions: local and global.

At the local level, the temporal layer encodes individual clauses or predicates that determine specific points on the timeline. These annotations highlight temporally salient events as presented in the text. The local temporal is not intended to be exhaustive. It rather marks those events whose temporal anchoring is clearly indicated by the author and may be important for interpretation. Therefore, its main function is to draw attention to temporally significant elements of the discourse, rather than to provide a complete temporal representation.

In the example in Figure 2, three such time points are annotated: t_1 – reference to Matilda’s existence, t_2 – Holmes referring to Matilda, t_3 – initiating a conversation about vampires.

At the global level, the temporal layer identifies broader time spans within which multiple events take place. These spans are typically unified by at least one aspect, e.g. shared participants or location. They group thematically coherent segments and provide a temporal frame for larger parts of the discourse. In this way, the global temporal layer structures the text by linking topic segments through

shared temporal contexts.

In the excerpt in Figure 2, we identify two main time spans: T_1 – the period of composing the letter, and T_2 – the conversation between Sherlock and Watson.

We assume that distinguishing a separate temporal layer is crucial for discourse analysis. Temporal information interacts with coherence relations but cannot be reduced to them. While coherence relations capture semantic connections between discourse units, temporal relations describe the ordering and duration of events in time. Combining these two types of relations would obscure their distinct roles: coherence relations organise meaning and interpretation, and temporal relations organise the timeline of events. By modelling temporality as a separate layer, the model preserves this distinction and allows for a more precise and flexible analysis of discourse structure.

Metatextual layer The metatextual layer, like the temporal layer, serves an organisational function in discourse, but it does so in a different way. The temporal layer structures discourse with respect to the ordering and duration of events and is therefore tied to its content. The metatextual layer, in turn, operates outside the represented events and organises the text as an artefact. It does not describe what happens, but rather how the text is presented, i.e., it captures elements that structure the text without contributing to its propositional content. This layer is the most explicit, as it is introduced by the

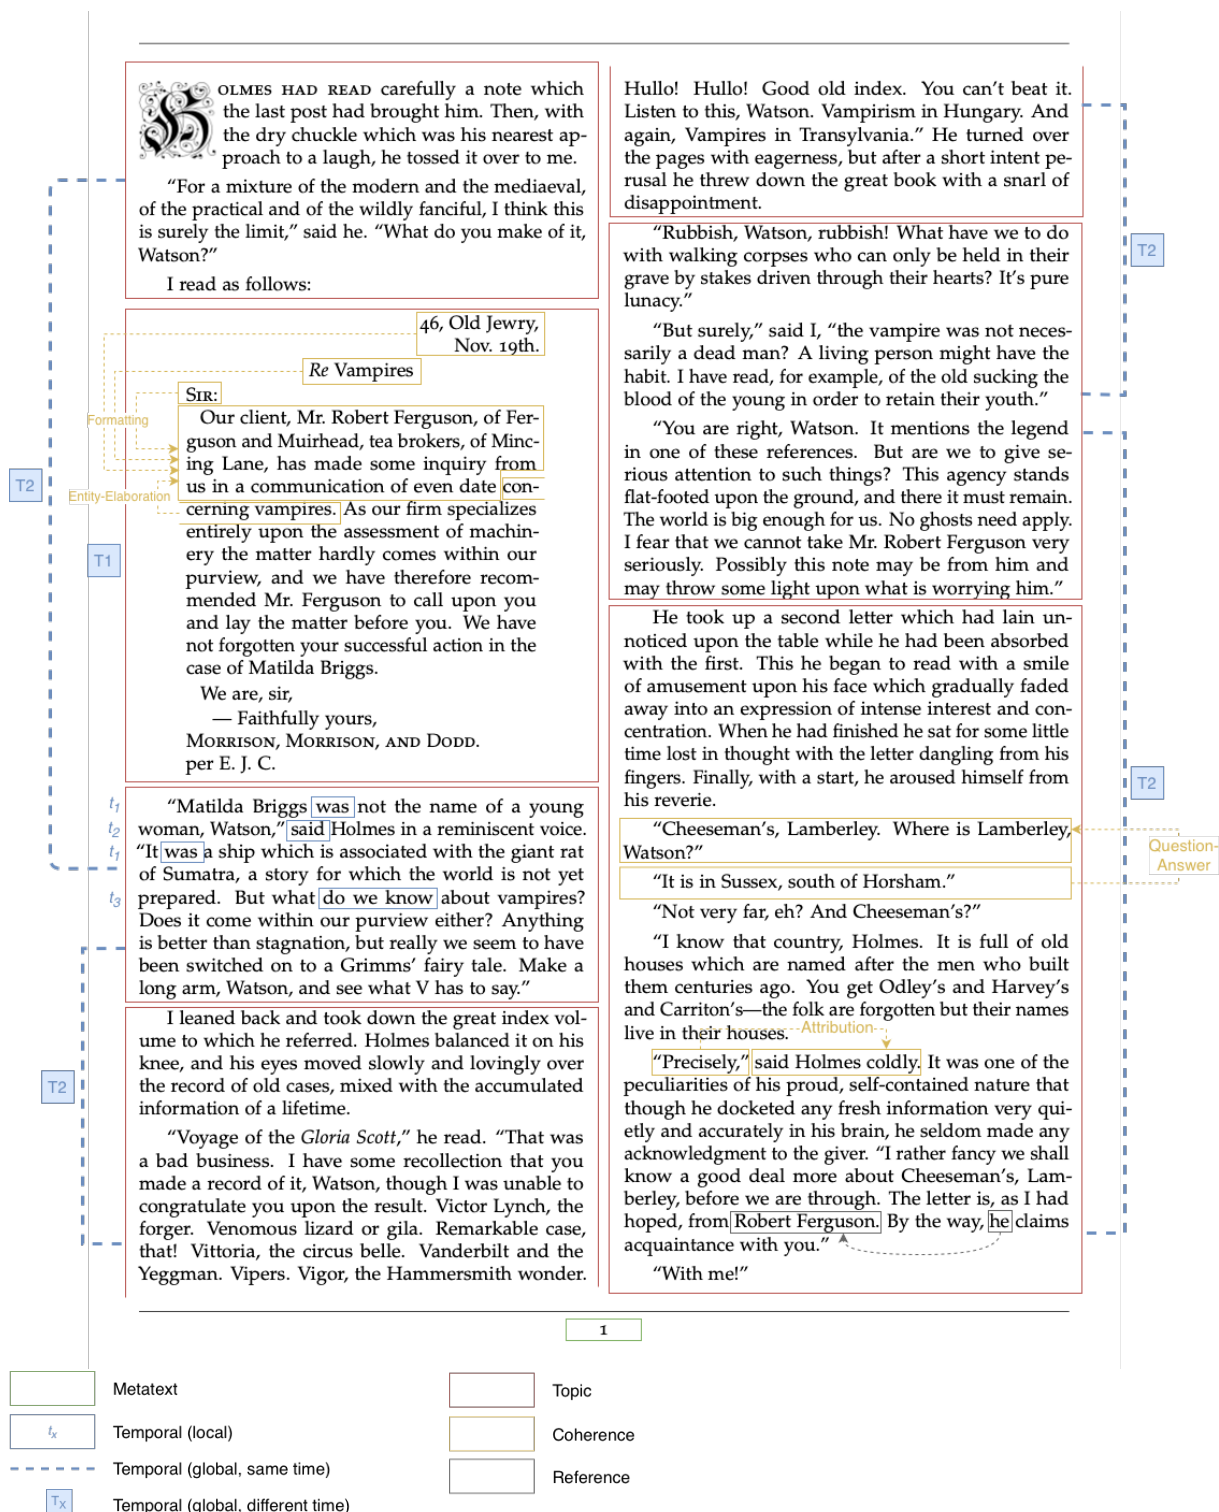


Figure 2: An excerpt from a Sherlock Holmes story (Doyle, 1927), annotated using the multilayer discourse model

author and directly reflects his/her intentions.

Metatextual segments include titles, headings, footnotes, page numbers (see Figure 2), and other editorial or formatting elements. They help organise the text, guide the recipient, and provide supplementary information. However, they do not contribute to the core meaning of the discourse in the same way as discourse segments do. For this reason, metatextual segments should not be linked to discourse segments by coherence relations. Their role is organisational or supplemental rather than semantic. We can illustrate this with the following example: even if a section title is replaced by a non-informative placeholder (e.g., ***), the reader can still understand the main body of the text. This shows that metatextual elements are not essential for interpreting discourse content.

By modelling the metatextual layer separately, our model distinguishes between content-related organisation (as in the temporal and topic layers) and external, author-imposed structuring.

Referential layer The referential layer captures relations between entities introduced in the text, mostly coreference relations. It models how entities are introduced and tracked across the discourse by linking expressions that refer to the same persons (see *Robert Ferguson* ← *he* in Figure 2). By connecting such expressions across discourse units, the referential layer helps identify discourse participants and interpret their roles.

Unlike coherence relations, referential relations are entity-based and operate at the level of entity tracking. They contribute to textual cohesion by maintaining referential continuity, but they do not determine the semantic structure of discourse. We include the referential layer to ensure that the proposed model provides a complete representation of the discourse.

Interaction between layers Although these layers capture different dimensions of discourse organisation, they do not operate independently. Instead, they interact and together shape discourse interpretation. The interaction between the topic layer and the coherence layer is particularly clear – coherence relations connecting discourse units establish topic segments (coherence relations do not extend across different topic segments). Topic segments may also show referential continuity. In narrative discourse, for example, a topic shift is often accompanied by a change in the participants involved in the described situation. Moreover, entities that will

later become referentially accessible are usually introduced explicitly at the beginning of a new topic segment. Topic shifts may also involve changes in temporal perspective, such as transitions between different time frames or events. Individual linguistic expressions may participate in multiple layers simultaneously. Metatextual markers, in turn, may explicitly signal topic transitions.

Distinguishing these layers analytically does not mean that discourse is strictly modular. Rather, the layered architecture is intended to capture partially autonomous organisational principles while also accounting for their systematic interaction in discourse interpretation and production.

4 Coherence Relations Decision Tree

According to what we presented in Section 3, in the future we aim to focus on a more comprehensive description of the document – one that accounts for multiple layers of coherence. For now, however, we have concentrated on the text itself and on the coherence relations that operate within it (in accordance with ISO). Consequently, our first step in developing the guidelines was to ensure annotation rigor by creating a decision tree.

The purpose of this approach is to organize coherence relations by: 1) assigning them to broader classes distinguished on the basis of the general types of interrelation that hold between two discourse segments and 2) providing a set of specified questions that help the annotator determine which type of relation is present in the text. Thanks to a hierarchical decision-making structure based on clear principles, we aim to achieve the highest possible level of unambiguity and to verify whether this approach allows us to increase annotation agreement.

We base our approach on characterizing the type of relation that holds between the segments by determining what function one segment performs with respect to the other, or what kind of meaning it conveys. For this reason, we aim to treat connectives linking discourse segments in a neutral way. We believe that relying on the presence of a particular connective can bias the judgment and lead to incorrect annotations.

Moreover, what we are trying to achieve is a universal description that – regardless of the linguistic material – can be applied to any language. A general comparison of one segment with respect to the other makes it possible to achieve this goal.

4.1 Tree structure

Navigation At the top of the tree lies the question: “What is the relation between X and Y?” There are four possible answers to this question: *Reliance*, *Broadening*, *Mutual*, and *Pragmatic*. Depending on which one is selected, the annotator proceeds to the corresponding branch, where further tests are located. Each question has been formulated in such a way that it can be answered with “yes” or “no”, allowing the annotator to determine whether the relation in question holds. If the answer is negative, we move on and respond to the next questions.

We should note here that although in the multilayer metamodel we distinguish *Attribution* and *Formatting* relations as ones that differ in nature from the core discourse relations (see Figure 1 and Section 3), in the decision tree we nevertheless include them among the relations that extend the narrative. Even though we treat them as a separate discourse level, during the annotation process we do not separate them, because, in very general terms, they spread the narrative, adding or specifying some information, cf. the first example of *Formatting* relations in Table 1 [*Karol Nowak wrote about this in his book*] [(15)]. Here, the second DU broadens the first DU, tightly linking to it, adding some reference. In the same time, the reference relates to another item, out of the proper discourse. To facilitate the annotators’ work, in such cases we grouped coherence relations from different layers into less fine-grained decision branches.

Motivation for the hierarchy Based on our work with the texts and the annotations we have carried out so far, we have established a hierarchy of relations. This does not mean that one relation is more important than another. Rather, it concerns the level of specificity of the relation that holds between two segments. For example, *Exemplification* specifies the situation presented in the first segment more precisely than *Elaboration* does. For this reason, the decision-making process begins with identifying the most specific relations, as well as those that are relatively the easiest to annotate (e.g., *Attribution* or *Formatting*, due to their strong grounding in the form of the analyzed text).

In the decision tree, each test (each question) is additionally accompanied by a set of maximally unambiguous examples. We present some of them in the Table 1 in the Appendix. This is why we do not repeat them in the subsections below.

Relations are based on the ISO set, although we

introduced several modifications:

- we merged the *Negative Condition* relation with *Condition*, assuming that both are grounded in the same underlying notion of conditionality;
- we excluded the *Expansion* relation from the set, treating it as a referential relation;
- we split *Elaboration* into two subtypes: *Elaboration* (where the second segment refers to the entire first segment) and *Entity-Elaboration* (where the second segment refers to a single element within the first segment). This distinction is substantial, and the *Elaboration* relation itself has a very broad scope that requires further specification;
- for the relations *Conjunction*, *Disjunction*, *Restatement*, *Functional dependence*, and *Feedback dependence*, we adopted the names *Co-occurrence*, *Alternative*, *Repetition*, *Question-Answer*, and *Feedback*, respectively, as we consider them more adequate and more accurately reflecting the nature of these relations;
- we added the *Comment* relation within the set of pragmatic relations.

Reliance relations All relations here follow a single general pattern: one segment stands in a relation of reliance to another segment. To identify a given relation, the annotator must answer the following questions:

Condition: *If X happens, does Y happen in consequence?*

Cause: *Does X cause Y? / Is X a reason for Y? / Is Y the result/effect of X?*

Concession: *Does Y occur despite X (on which a different result could depend)?*

Purpose: *Does Y indicate the goal of X?*

Substitution: *Does Y replace X?*

Broadening relations Broadening relations expand, specify, or refine the first discourse unit by the second one, adding new information:

Formatting: *Is Y relative to X a footnote, reference or a list item?*

Attribution: *Does Y identify the source or author of the information/statement expressed in X?*

Manner: *Does Y indicate the manner in which X occurs?*

Exemplification: Does *Y* provide an example of *X*? *X* is a set of e.g. features/items and *Y* is one of those features/items

Exception: Does *Y* state an exception to the rule given in *X* (being part of *X* but excluded from it)?

Elaboration: Does *Y* expand or specify the situation presented in the whole of *X*, adding some new information?

Entity-Elaboration: Does *Y* expand or specify one element presented in *X*?

Mutual relations These are symmetric relations – the two discourse segments can be interchanged without altering the overall interpretive meaning:

Repetition (Restatement): Does *X* mean the same as *Y* (expressed in the same or different words)?

Similarity: Is *Y* similar/comparable to *X*?

Alternative (Disjunction): Can either *X* or *Y* occur (mutually exclusive) and is there any choice between *X* or *Y*?

Contrast: Does *X* contrast with *Y*? / Is *X* opposed to *Y*?

Co-occurrence (Conjunction): Do *X* and *Y* co-occur in the discourse?

Pragmatic relations The distinctive feature of these relations is that their primary function is pragmatic. They have the character of dialogic acts or pragmatic comments, in which the interaction between the two discourse segments comes to the foreground:

Question–Answer (Functional dependence): Is *Y* an answer to the question raised in *X*?

Feedback (Feedback dependence): Is *Y* a feedback response to *X* other than a question?

Comment: Is *Y* a comment on *X*?

4.2 Annotation order

In the first stage of annotation, before we enter the decision tree, we identify relations whose nature is clearly different from those that are of primary interest to us: metatextual, topical, temporal, and referential. Next, we focus on the coherence relations themselves. Then, we follow the scheme (the same for all languages) described in Section 4.1. Identifying each relation requires returning to the top of the tree every time.

In annotation, it is also important to name segments in asymmetric relations in a way that preserves the dependency structure. ISO defines an asymmetric relation as follows: (REL, A, B) is not equivalent to (REL, B, A). Wherever possible, we

follow ISO naming conventions (ISO, 2016, pp. 13–14). Where this is not possible, we introduce our own labels (e.g., in the case of *Attribution* X = *Attributor* and Y = *Source*).

The goal of the annotation is to cover the entire text with relations that will ultimately form a relation tree. We use an annotation tool that we are developing ourselves specifically for the project, and we continue to improve and modify it according to our current needs.

5 Conclusions

In the first phase of our work (Ogrodniczuk et al., 2025) we focused on segmenting the text into discourse units. Here we have shown how we plan to link these units with coherence relations in a comprehensive framework for the 'universal' approach to discourse annotation, successfully synthesizing the established ISO 24617-8 semantic annotation framework with a newly proposed multi-layer model of coherence relations. By separating discourse into distinct, yet interacting layers – such as the coherence, metatextual, temporal, and referential layers – we try to distinguish the unique functional roles associated with different linguistic phenomena.

Our proposed decision tree operationalizes this theoretical model, offering a systematic, question-driven guide that effectively minimizes subjective bias and maximizes inter-annotator agreement. The flat representation that emerges from this process retains the practical benefits of previous frameworks while ensuring a higher degree of cross-framework interoperability.

Moving forward, we plan to apply this universal framework to a wider variety of corpora and diverse languages to validate its robustness and adaptability.

Acknowledgments

This research was funded in whole by the National Science Centre, Poland, grant 2023/50/A/HS2/00559 ("Universal Discourse: a multilingual model of discourse relations").

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018744.

References

- Nicholas Asher and Alex Lascarides. 1993. [Temporal Interpretation, Discourse Relations and Commonsense Entailment](#). *Linguistics and Philosophy*, 16:437–493.
- Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge.
- Nicholas Asher and Laure Vieu. 2005. [Subordinating and Coordinating Discourse Relations](#). *Lingua*, 115(4):591–610.
- Doug Beeferman, Adam L. Berger, and John D. Lafferty. 1999. [Statistical Models for Text Segmentation](#). *Machine Learning*, 34:177–210.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet Allocation](#). *Journal of Machine Learning Research*, 3:993–1022.
- Chloé Braud, Amir Zeldes, Chuyuan Li, Yang Janet Liu, and Philippe Muller. 2025. [The DISRPT 2025 Shared Task on Elementary Discourse Unit Segmentation, Connective Detection, and Relation Classification](#). In *Proceedings of the 4th Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2025)*, pages 1–20, Suzhou, China. Association for Computational Linguistics.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski. 2001. [Building a Discourse-Tagged Corpus in the Framework of Rhetorical Structure Theory](#). In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*, Aalborg, Denmark. Association for Computational Linguistics.
- Liesbeth Degand and Jacqueline Evers-Vermeul. 2015. [Grammaticalization or pragmaticalization of discourse markers?: More than a terminological issue](#). *Journal of Historical Pragmatics*, 16(1):59–85.
- Arthur Conan Doyle. 1927. [The Adventure of the Sussex Vampire](#). In *The Casebook of Sherlock Holmes*. John Murray. First published in The Strand Magazine in 1924.
- Herbert Paul Grice. 1975. Logic and Conversation. *Syntax and Semantics*, 3:41–58.
- Michael A. K. Halliday. 2014. *An Introduction to Functional Grammar*, 4th edition. Routledge, London and New York.
- Michael A. K. Halliday and Ruqaiya Hasan. 1976. *Coherence in English*. Longman, London.
- Marti A. Hearst. 1997. [Text tiling: Segmenting text into multi-paragraph subtopic passages](#). *Computational Linguistics*, 23(1):33–64.
- Jerry R. Hobbs. 1979. [Coherence and Coreference](#). *Cognitive Science*, 3(1):67–90.
- ISO. 2014. [ISO 24617-5:2014: Language resource management — Semantic annotation framework \(SemAF\) — Part 5: Discourse structure \(SemAF-DS\)](#). Standard ISO 24617-5:2014, International Organization for Standardization, Geneva, Switzerland.
- ISO. 2016. [ISO 24617-8:2016: Language resource management – Semantic annotation framework \(SemAF\) – Part 8: Semantic relations in discourse \(DR-Core\)](#). Standard ISO 24617-8:2016, International Organization for Standardization, Geneva, Switzerland.
- ISO. 2020. [ISO 24617-2:2020: Language resource management — Semantic annotation framework \(SemAF\) — Part 2: Dialogue acts](#). Standard ISO 24617-2:2020, International Organization for Standardization, Geneva, Switzerland.
- Daniel Jurafsky and James H. Martin. 2026. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, with Language Models*, 3rd edition. Online manuscript released January 6, 2026.
- Walter Kintsch. 1998. *Comprehension: A Paradigm for Cognition*. Cambridge University Press, Cambridge.
- Alex Lascarides and Nicholas Asher. 2008. [Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure](#). In *Computing Meaning. Studies in Linguistics and Philosophy*, volume 83, pages 87–124. Springer, Dordrecht, Netherlands.
- William C. Mann and Sandra A. Thompson. 1987. [Rhetorical Structure Theory: A Theory of Text Organization](#). Technical report, Information Sciences Institute, Marina del Rey, CA.
- Maciej Ogrodniczuk, Anna Latusek, Karolina Saputa, Alina Wróblewska, Daniel Ziembicki, Bartosz Żuk, Martyna Lewandowska, Adam Okrański, Paulina Rosalska, Anna Śliwicka, Aleksandra Tomaszewska, and Sebastian Żurowski. 2025. [Where frameworks \(dis\)agree: A study of discourse segmentation](#). In *Proceedings of the 6th Workshop on Computational Approaches to Discourse, Context and Document-Level Inferences (CODI 2025)*, pages 182–196, Suzhou, China. Association for Computational Linguistics.
- Silvia Pareti. 2016. [PARC 3.0: A Corpus of Attribution Relations](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3914–3920, Portorož, Slovenia. European Language Resources Association (ELRA).
- Volha Petukhova, Laurent Prévot, and Harry Bunt. 2011. [Multi-level discourse relations between dialogue units](#). In *Proceedings of the Sixth Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation (ISA-6)*, pages 18–27, Oxford, United Kingdom. Association for Computational Linguistics.

- Andrew Potter. 2019. [The Rhetorical Structure of Attribution](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 38–49, Minneapolis, MN. Association for Computational Linguistics.
- 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'08)*, pages 2961–2968, Marrakech, Morocco. European Language Resources Association (ELRA).
- Rashmi Prasad, Bonnie Webber, and Aravind Joshi. 2014. [Reflections on the Penn Discourse TreeBank, Comparable Corpora, and Complementary Annotation](#). *Computational Linguistics*, 40(4):921–950.
- James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Rob Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir Radev. 2003. [TimeML: Robust Specification of Event and Temporal Expressions in Text](#). pages 28–34.
- Manfred Stede and Arne Neumann. 2014. [Potsdam commentary corpus 2.0: Annotation for discourse research](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 925–929, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Teun A. van Dijk. 1980. *Macrostructures: An Interdisciplinary Study of Global Structures in Discourse, Interaction, and Cognition*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Bonnie Webber, Markus Egg, and Valia Kordoni. 2012. [Discourse Structure and Language Technology](#). *Natural Language Engineering*, 18(4):437–490.
- Florian Wolf, Edward Gibson, Amy Fisher, and Meredith Knight. 2005. [Discourse Graphbank](#).
- Sandrine Zufferey and Liesbeth Degand. 2024. [Connectives and Discourse Relations](#). Key Topics in Semantics and Pragmatics. Cambridge University Press, Cambridge.
- Urszula Żydek-Bednarczuk. 2005. *Wprowadzenie do lingwistycznej analizy tekstu*. Universitas, Kraków.

What is the relation between X and Y?

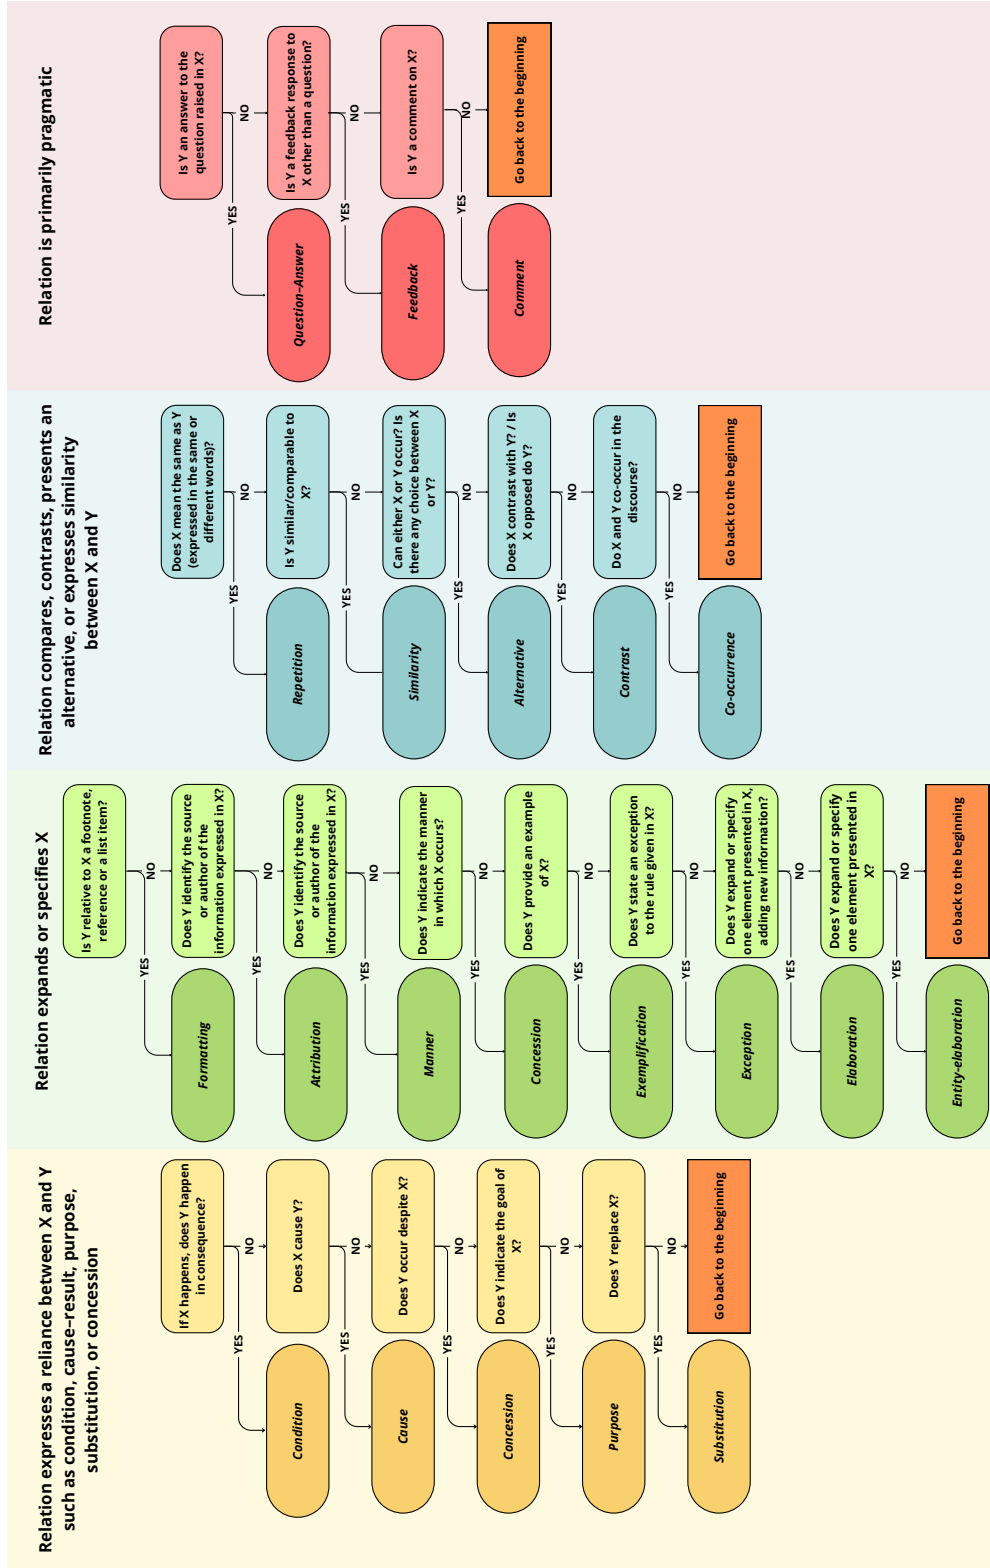


Figure 3: Annotation decision tree

Relation Type	Definition / Paraphrase	English Example
Reliance Relations		
Condition	If X happens, then Y happens	[You will break your leg] [if you keep jumping.] [If I hadn't been jumping,] [I wouldn't have broken my leg.]
Cause	X causes Y / Y results from X	[Sophia broke her leg] [because she was jumping.] [I don't smoke cigarettes] [so I don't pollute the air.]
Concession	Y happens despite X	[Although I was in a great hurry] [I didn't manage to buy bread.] [They lost the match] [even though they gave it their all.]
Purpose	Y indicates the goal of X	[I am studying English] [so that I can study at Oxford.]
Substitution	Y replaces X	[Instead of going to bed] [he read a book until midnight.]
Broadening Relations		
Manner	Y indicates the way in which X occurs	[I resolved the matter in such a way] [that I simply let it go.]
Exemplification	Y provides an example of X	[Steel parts of American towers ended up in eleven countries.] [They were sent, for example, to India, China, and Malaysia.] [We have plenty of it.] (books) [Let me start with the classic of classics, namely Star Trek.]
Exception	Y is an exception to the rule given in X	[All students failed the exam] [except Iza, who got the maximum score.]
Elaboration	Y expands or specifies the situation presented in the whole of X, adding some new information	[Maria received a prestigious scientific award] [– namely a grant from the European Research Council.]
Entity-Elaboration	Y expands or specifies one element in X	[There was a small distance between the gates and the tunnel] [which the runners finishing the race entered.] [Ewa bought a new car.] [The car is red.]
Mutual Relations		
Repetition (Restatement)	Y restates or repeats X	[He is very busy] [that is, he doesn't have time now.]
Similarity	Y is similar or comparable to X	[He spoke passionately] [just as his father used to do years ago.]
Contrast	Y contrasts with X	[I like winter] [and she likes summer.]
Alternative (Disjunction)	X and Y are mutually exclusive; It can be a choice between X and Y	[In the evening I will watch a film] [or I will read a book.]
Co-occurrence (Conjunction)	X and Y co-occur in the discourse	[I don't pollute the air] [and I don't litter.]
Pragmatic Relations		
Question–Answer (Functional Dependence)	Y answers the question raised in X	[Did I really do the right thing?] [I think so.]
Feedback (Feedback Dependence)	Y is a feedback response to X (not a question)	[– This pasta is good.] [– Mhm.] [– I will borrow your car this evening.] [– OK.]
Comment	Y is a comment to X	[– Frankly speaking,] [this chicken is very tasty.]
Attribution Relations		
Attribution	Y is attributed to a speaker or source expressed in X	[I like going to work,] [said Peter.] [The constitution says:] [(constitutional article)]
Formatting Relations		
Formatting	Y is a footnote, reference, or list item relative to X	[Karol Nowak wrote about this in his book] [(15).] [A long time ago there were no telephones] [(see “Telephones and Modernity” by Nowak).] [To obtain an ID card, one must:] [– provide a copy of a photo,] [– present proof of payment.]

Table 1: Coherence relations