
**Language resource management —
Semantic annotation framework
(SemAF) —**

**Part 8:
Semantic relations in discourse, core
annotation schema (DR-core)**

*Gestion des ressources langagières — Cadre d'annotation sémantique
(SemAF) —*

*Partie 8: Relations sémantiques dans le discours, schéma d'annotation
de base (DR-core)*

STANDARDSISO.COM : Click to view the full PDF of ISO 24617-8:2016



COPYRIGHT PROTECTED DOCUMENT

© ISO 2016, Published in Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Ch. de Blandonnet 8 • CP 401
CH-1214 Vernier, Geneva, Switzerland
Tel. +41 22 749 01 11
Fax +41 22 749 09 47
copyright@iso.org
www.iso.org

Contents

Page

Foreword	iv
Introduction	v
1 Scope	1
2 Normative references	2
3 Terms and definitions	2
4 Basic concepts and metamodel	3
4.1 Overview	3
4.2 Representation of discourse structure	3
4.3 Semantic description of discourse relations	4
4.4 Pragmatic variants of discourse relations	4
4.5 Hierarchical classification of discourse relations	5
4.6 Inference of multiple relations between two segments	5
4.7 Representation of (a)symmetry of relations	6
4.8 Representation of the relative importance of arguments for discourse meaning/ structure	6
4.9 Arity of arguments	7
4.10 Syntactic form, extent, and (non-)adjacency of argument realizations	7
4.11 Triggers of discourse relations	7
4.12 Representation of attribution as a discourse relation	8
4.13 Representation of entity-based relations	9
4.14 Representation of non-existence of a discourse relation	10
4.15 Summary: Assumptions of the DR-core annotation scheme	10
4.16 Issues to be taken up in the follow-up of DR-core	11
4.17 Metamodel	11
5 Core discourse relations	12
6 Current approaches and annotation schemes	21
6.1 Overview	21
6.2 Rhetorical structure theory (RST)	21
6.3 RST Treebank	22
6.4 Hobbs' Theory of Discourse Coherence (HTDC)	24
6.5 GraphBank	24
6.6 SDRT	25
6.7 CCR	26
6.8 Penn Discourse Treebank (PDTB)	26
6.9 Mapping of DR-core discourse relations to existing classifications	28
7 Interactions of this document with other annotation schemes	30
7.1 Overlapping annotation schemes	30
7.2 Discourse relations and semantic roles	31
7.3 Discourse relations and temporal relations	31
7.4 Discourse relations and semantic relations between dialogue acts	32
8 DRelML: Discourse Relations Markup Language	33
8.1 Overview	33
8.2 DRelML abstract syntax and semantics	34
8.3 Concrete syntax	35
Bibliography	39

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see the following URL: www.iso.org/iso/foreword.html.

The committee responsible for this document is ISO/TC 37, *Terminology and other language and content resources*, Subcommittee SC 4, *Language resource management*.

A list of all parts in the ISO 24617 series can be found on the ISO website.

Introduction

The last decade has seen a proliferation of linguistically annotated corpora coding many phenomena in support of empirical natural language research, both computational and theoretical. At the level of discourse, interest in discourse processing has led to the development of several corpora annotated for discourse relations. Discourse relations, also called “coherence relations” or “rhetorical relations”, are relations, expressed explicitly or implicitly, between situations mentioned in a discourse and are key to a complete understanding of the discourse, beyond the meaning conveyed by clauses and sentences. Discourse relations and discourse structure are considered to be key ingredients for NLP tasks such as summarization,^{[39][41]} complex question answering,^[74] natural language generation,^{[19][47][56]} machine translation,^[42] opinion mining and sentiment analysis,^{[11][12]} and information retrieval.^[38] A recent overview^[76] includes a description of the state of the art in discourse and computation. Several international and collaborative efforts have resulted in annotated resources of discourse relations, across languages as well as genres, to support the development of such applications.

Existing annotation frameworks exhibit two major differences in their underlying assumptions, one of which concerns the representation of discourse structure, while the other has to do with the semantic classification of discourse relations. As a result, annotations constructed using one framework are not easily interpreted in another framework, and annotated resources are limited in their interoperability. Notwithstanding their differences, however, there are strong compatibilities between them that can be clarified and used as the basis for mappings and comparisons between the resources, as well as for use as a basis for future annotation.

In a coherent (written or spoken) discourse, the situations mentioned in the discourse, such as events, states, facts, propositions, and dialogue acts are semantically linked through causal, contrastive, temporal and other relations, called “discourse relations”, “rhetorical relations”, or “coherence relations”. Although discourse relations hold most prominently between the meanings of successive sentences or utterances in a discourse, they may also occur between the meanings of smaller or larger units (nominalizations, clauses, paragraphs, dialogue segments), and they may occur between situations that are not explicitly described but that can be inferred.

This document aims to specify an interoperable approach to the annotation of local semantic relations in discourse (DRels), following the Linguistic Annotation Framework (LAF, ISO 24612-2; see also Reference ^[23]) and the general principles for semantic annotation established in ISO 24617-6. It reflects the view that strong underlying compatibilities with respect to the semantic description of discourse relations can be observed in the various discourse relation frameworks being used to support data annotation, e.g. Rhetorical Structure Theory (RST),^[40] Segmented Discourse Representation Theory (SDRT),^[3] the Penn Discourse Treebank,^[59] Hobbs’ Theory of Discourse Coherence (HTDC)^{[17][18]} and the Cognitive Approach to Coherence Relations (CCR)^[66]. This document aims to provide an explanation of these compatibilities and a loose mapping between definitions of individual discourse relations, as specified in the different frameworks that will benefit the community as a whole.

The main aims of this document are to (1) establish a set of desiderata for interoperable DRel annotation; (2) specify a way of annotating DRels that is compatible with existing and emerging ISO standard annotation schemes for semantic information; and (3) provide clear and mutually consistent definitions of a set of “core” discourse relations which are commonly found in some form in many existing discourse relation frameworks. Together, (2) and (3) form a “core annotation scheme” for DRels.

This document does not aim at providing a fixed and exhaustive set of discourse relations, but rather at providing an open, extensible set of core relations. The core annotation scheme also discusses certain issues in discourse relation annotation that it leaves open, as they require further study in collaboration with other efforts in multilingual discourse annotation, in particular the European COST action TextLink. A future part of ISO 24617 is envisaged that will complement this document by providing a complete interoperable annotation scheme for DRels, while also addressing the multilingual dimension of the standard. The issues to be taken up for this complementary part are listed in [4.16](#).

STANDARDSISO.COM : Click to view the full PDF of ISO 24617-8:2016

Language resource management — Semantic annotation framework (SemAF) —

Part 8:

Semantic relations in discourse, core annotation schema (DR-core)

1 Scope

This document establishes the representation and annotation of local, “low-level” discourse relations between situations mentioned in discourse, where each relation is annotated independently of other relations in the same discourse.

This document provides a basis for annotating discourse relations by specifying a set of core discourse relations, many of which have similar definitions in different frameworks. To the extent possible, this document provides mappings of the semantics across the different frameworks.

This document is applicable to two different situations:

- for annotating discourse relations in natural language corpora;
- as a target representation of automatic methods for shallow discourse parsing, for summarization, and for other applications.

The objectives of this specification are to provide:

- a reference set of data categories that define a collection of discourse relation types with an explicit semantics;
- a pivot representation based on a framework for defining discourse relations that can facilitate mapping between different frameworks;
- a basis for developing guidelines for creating new resources that will be immediately interoperable with pre-existing resources.

With respect to discourse structure, the limitation of this document to specifications for annotating local, “low-level” discourse relations is based on the view that (a) the analysis at this level is what is well understood and can be clearly defined; (b) further extensions to represent higher-level, global discourse structure is possible where desired; and (c) that it allows for the resulting annotations to be compatible across frameworks, even when they are based on different theories of discourse structure.

As a part of the ISO 24617 semantic annotation framework (“SemAF”), the present DR-core standard aims to be transparent in its relation to existing frameworks for discourse relation annotation, but also to be compatible with other ISO 24617 parts. Some discourse relations are specific to interactive discourse, and give rise to an overlap with ISO 24617 Part 2, the ISO standard for dialogue act annotation. Other discourse relations relate to time, and their annotation forms part of ISO 24617-1 (time and events); still other discourse relations are very similar to certain predicate-argument relations (“semantic roles”), whose annotation is the subject matter of ISO 24617-4. Since the various parts are required to form a consistent whole, this document pays special attention to the interactions of discourse relation annotation and other semantic annotation schemes (see [Clause 8](#)).

This document does not consider global, higher-level discourse structure representation which involves linking local discourse relations to form one or more composite global structures.

This document is, moreover, restricted to strictly *semantic* relations, to the exclusion of, for example, presentational relations, which concern the way in which a text is presented to its readers or the way in which speakers structure their contributions in a spoken dialogue.

2 Normative references

There are no normative references in this document.

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1 discourse

sequence of clauses or sentences in written text or of utterances in oral speech

3.2 situation

eventuality, fact, proposition, condition, belief or dialogue act, that can be realized by a linguistically simple or complex expression, such as a clause, a nominalization, a sentence/utterance, or a discourse segment consisting of multiple sentences or utterances

3.3 discourse relation

relation between two *situations* (3.2) mentioned in a *discourse* (3.1)

EXAMPLE 1 “Peter came late to the meeting. He had been in a traffic jam.” The events mentioned in the two sentences are implicitly related through the discourse relation *Cause*.

EXAMPLE 2 “Peter was in a traffic jam, but he arrived on time for the meeting.” The events mentioned in the two clauses are related by the discourse relation *Concession*, expressed by the connective “but”.

EXAMPLE 3 “Peter did not manage to come to the meeting; he was held up in a terrible traffic jam.” The causal relation in this example is the same as in Example 1, but the argument expressed by the first clause is not an eventuality, but a proposition, formed by an event description with negative polarity.

Note 1 to entry: Quasi-synonyms for “discourse relation”, with small variations in meaning, are “coherence relation” and “rhetorical relation”.

3.4 discourse connective

word or multi-word expression expressing a *discourse relation* (3.3)

EXAMPLE Single-word discourse connectives include “but”, “since”, “and”, “however”, “because”. Multi-word discourse connectives include “as well as”, “such as”.

Note 1 to entry: Many of the words that can be used as discourse connectives can also be used as intra-clausal conjunctions, as with the use of “and” in “John and Mary are a lovely couple”.

3.5 low-level discourse structure

representation of discourse structure that only specifies local dependencies between a discourse relation and its arguments, without further specifying any links or dependencies across these local structures

4 Basic concepts and metamodel

4.1 Overview

In a discourse, which comes into play when communication involves a sequence of clauses or sentences in a text, or utterances in a dialogue, a major aspect of the understanding comes from how the events, states, facts, propositions, and dialogue acts mentioned in the discourse are related to each other. Understanding such relations, such as Cause, Contrast, and Condition, contribute to what is called the “coherence” of the discourse, and they can be “realized” explicitly, by means of certain words and phrases (often called “connectives”), or they can be implicit, when they have to be inferred on the basis of the discourse context and world knowledge. Examples 1 to 3 illustrate the Cause relation realized with expressions from different syntactic classes. In Example 1, a subordinating conjunction “because” is used to connect some situation (here, the meaning of the subordinate clause) as the reason for the buying event mentioned in its matrix clause. In Example 2, an adverb “as a result” is used to relate two sentences to express the consequence of not seeing many signs about growth coming to a halt. In Example 3, an explicit phrase is again used, to explain the claim about the level of investor withdrawal, but here the phrase does not correspond to a well-defined single syntactic class such as a conjunction or adverb. Finally, Example 4 shows that although a causal relation can be inferred between the two sentences, with the second sentence offering an explanation for why some (investors) have raised their cash positions, there is no word or phrase in the text to express this inference. Rather, the discourse context needs to be used together with, cohesive devices and world knowledge to get at the relation. Often, when such relations are inferred, it is possible to insert a connective phrase^[44] to express the relation, as shown here with the insertion of “because”. In this document, the term “connective” is used in a broad sense, to refer to any word or phrase used to express a discourse relation, including both those drawn from well-defined syntactic classes as well as those that are not.

Example 1 Mr. Taft, who is also president of Taft Broadcasting Co., said *he bought the shares because he keeps a utility account at the brokerage firm of Salomon Brothers Inc., which had recommended the stock as a good buy.*

Example 2 *Despite the economic slowdown, there are few clear signs that growth is coming to a halt. As a result, Fed officials may be divided over whether to ease credit.*

Example 3 *But a strong level of investor withdrawal is much more unlikely this time around, fund managers said. A major reason is that investors already have sharply scaled back their purchases of stock funds since Black Monday.*

Example 4 *Some have raised their cash positions to record levels. [implicit (because)] High cash positions help buffer a fund when the market falls.*

Existing frameworks for describing and representing discourse relations differ along several lines. The remainder of this clause provides a comparison of the most important frameworks, focusing on those that have been used as the basis for annotating discourse relations in corpora, in particular the Theory of Discourse Coherence (HTDC)¹⁾ by Hobbs,^[18] Rhetorical Structure Theory (RST) by Mann and Thompson,^[40] the Cognitive Approach of Coherence Relations (CCR) by Sanders and others,^[66] Segmented Discourse Representation Theory (SDRT) by Asher and Lascarides^[3] and the annotation framework of the Penn Discourse Treebank (PDTB).^{[59][61]} The comparison highlights and discusses the main issues that are considered relevant for developing the pivot representation in DR-core. For each issue, the discussion is followed by the ISO specification adopted for that issue. The clause ends with a summary of the key features of the DR-core specification, and the DR-core metamodel.

4.2 Representation of discourse structure

One important difference between existing DRel frameworks concerns the representation of discourse structure. For example, the RST Treebank,^[10] based on the Rhetorical Structure Theory,^[40] assumes a tree representation to subsume the entire text of the discourse. The Discourse GraphBank,^[78] based on

1) “HTDC” as an acronym for Hobbs’ theory is created for the purpose of this document and does not, thus far, appear elsewhere in the literature.

HTDC, allows for general graphs that permit multiple parents and crossing, and the DISCOR corpus^[64] and the ANNODIS corpus,^[1] based on SDRT, allow directed acyclic graphs that permit multiple parents, but not crossing. There are also frameworks that are pre-theoretical or theory-neutral with respect to discourse structure. These include the PDTB,^[59] based loosely on a lexicalized approach to discourse relations and structure (DLTAG^{[16][75]}, and DiscAn,^[65] based on CCR). In both of these frameworks, individual relations along with their arguments are annotated, without being combined with other relations to form a composite structure encompassing the entire text.

These widely different views about the structural representation for discourse are difficult to reconcile with each other. *In the DR-core specification, a pre-theoretical stance involving low-level annotation of discourse relations is adopted*, with the idea that individual relations can be more reliably annotated and that they can be further annotated to project a higher-level tree or graph structure, depending on one's theoretical inclination. From the point of view of interoperability, the low-level annotation can also serve as a pivot representation when comparing annotations of different resources grounded in different theories.

4.3 Semantic description of discourse relations

A second difference among existing frameworks relates to whether the meaning of a discourse relation is described in “informational” term, i.e. in terms of the “meaning” of the relation's arguments, or in “intentional” terms, i.e. in terms of the intentions of the speaker/writer (W) and intended effects on the hearer/reader (R). While SDRT, HTDC, PDTB and CCR describe the meaning in informational terms, RST provides definitions in intentional terms. For instance, Example 5 shows the definition for the (non-volitional) Cause relation in RST (N = nucleus, S = satellite, W = writer, R = reader), while Example 6 presents the definition for the same relation in HTDC (where it is called Explanation).

Example 5 *Non-Volitional Cause (RST)*

Constraints on N: presents a situation that is not a nucleus

Constraints on the N + S combination: S presents a situation that, by means other than motivating a volitional action, caused the situation presented in N; without the presentation of S, R might not know the particular cause of the situation; a presentation of N is more central than S to W's purposes in putting forth the N-S combination

The effect: R recognizes the situation presented in S as a cause of the situation presented in N

Locus of the effect: N and S.

Example 6 *Explanation (HTDC)*

Infer that the state/event asserted by S_1 causes or could cause the state/event asserted by S_0 .

Despite the different ways of describing DRel semantics, it is important to note that in many cases, the differences lie in the “level” at which the relation is described, especially when the situations being related are the same. Thus, for example, a DRel defined in informational terms in one framework can be effectively mapped to a DRel in another framework where it may be defined in intentional terms. With this in mind, *DRel meaning in the DR-core specification is described in “informational” terms*, but in 6.9, a mapping is provided from the core relation types (presented in [Clause 5](#)) to the relations present in existing classifications, including those that define relations in intentional terms.

4.4 Pragmatic variants of discourse relations

With the exception of HTDC, all frameworks also distinguish relations when one or both of the arguments involve an implicit belief or a dialogue act²⁾ that takes scope over the semantic content of the argument. The motivation for this distinction comes from examples like Example 7, where it should *not* be inferred that John's sending of the message somehow led to him being absent from work, but

2) The concept of a dialogue act, as used in ISO 24617-2, can be seen as an empirically based and computationally well-defined interpretation of the traditional notion of a “speech act”.

rather that it causes the speaker/writer to believe that John is not at work. In other words, the meaning of the subordinate clause provides evidence supporting the claim made by the main clause. Similarly, in Example 8, the inference should be made that the explanation is being provided not for the content of the question but for the (dialogue) act of questioning itself.

Example 7 *John is not at work today, because he sent me a message to say he was sick.*

Example 8 *What are you doing tonight? Because there's a good movie on.*

This kind of distinction has been given various names in the literature, for example the “semantic-pragmatic” distinction,[73][66][46] the “internal-external” distinction,[17][44] the “ideational-pragmatic” distinction[63] and the “content-metatalk” distinction.[37] In other cases, such as in RST, the distinction, while not being explicitly named, is evidently taken into account in the classification (e.g. Cause vs. Evidence/Justify in RST distinguishes the semantic and pragmatic interpretations, respectively). What is difficult to reconcile about the treatment of this distinction across the various frameworks is that while some, like CCR, allow for it for all relation types, others, like the PDTB and RST, only admit it for some relations (e.g. Cause, Condition, Contrast, Concession in PDTB). It must be noted, however, that there doesn't seem to be any a priori reason for such a restriction to only some relation types, and the choice is in the end found to result from what was observed in the corpus that was analysed and/or annotated. *In DR-core, the “semantic-pragmatic” distinction is allowed for all relation types, with the general aim of not being overly restrictive in the absence of well-defined criteria. At the same time, the scheme does not encode this distinction on the relation, but rather on the arguments of the relation, the main reason being that in all cases involving either a belief or a dialogue act, what is different is not the relation, but rather the semantic status of the arguments. A further motivation comes from recognizing that representing the distinction on the relation would not distinguish cases where the belief or dialogue act is implicit (as in Examples 7 and 8) from those where they are made explicit with performative verbs or propositional attitude verbs, as in Examples 9 and 10. Pragmatic interpretations are therefore represented on arguments using a feature indicating the argument to be of the type “belief” or the type “dialogue act”. Note that in cases exemplified by Examples 9 and 10 the belief or dialogue act aspect of the meaning is entirely obtained from the explicit content of the arguments, rather than from a contextually motivated inference.*

Example 9 *I believe John is not at work today because he sent me a message to say he was sick.*

Example 10 *I'm asking you what you are doing tonight because there's a good movie on.*

4.5 Hierarchical classification of discourse relations

In all existing frameworks, discourse relations are grouped together semantically to a greater or lesser degree; where they differ is in how the groupings are done. For example, while PDTB groups Concession together with Contrast under the broader Comparison class, CCR places Concession under the Negative Causal relation group, while placing Contrast under the Negative Additive group. Reconciliation with respect to these groupings is not possible, since they stem from basic differences in what is taken to count as semantic closeness. *The solution adopted in the DR-core specification is to use a “flat” set of core relations that can be used in an annotation scheme as just that, or mapped to the appropriate type within a particular hierarchical scheme adopted. In 6.9, these mappings from the DR-core relations to the schemes in different frameworks are provided.*

4.6 Inference of multiple relations between two segments

Among the various frameworks, the PDTB is unique in allowing multiple relations to be inferred between two given situations. The connective “since”, for example, can have both temporal and causal interpretations, as in Example 11.

Example 11 *MiniScribe has been on the rocks since it disclosed earlier this year that its earnings reports for 1988 weren't accurate.*

The DR-core specification provides for representing multiple relations inferred between two given situations, both when the relations are realized explicitly as well as implicitly.

4.7 Representation of (a)symmetry of relations

Whether or not a discourse relation is symmetric or asymmetric is a distinction embodied in the representation of all frameworks. That is, given a relation REL and its arguments A and B, all frameworks distinguish whether or not (REL, A, B) is equivalent to (REL, B, A). For example, the Contrast relation is taken to be symmetric whereas the Cause relation is considered asymmetric. Where frameworks differ is in how this distinction is captured in the scheme. Most classifications, such as RST, CCR, HTDC and PDTB, encode asymmetry in terms of the textual linear ordering and/or the syntax of the argument realizations. Thus, in the CCR classification, where the argument span ordering is one of the basic “cognitive” primitives underlying the scheme, the relation Cause-Consequence captures the “basic” order for the semantic causal relation, with the cause appearing before the effect, whereas the relation Consequence-Cause captures the “non-basic” order, with the effect appearing before the cause. In the PDTB, argument spans are first named as Arg1 and Arg2 according to syntactic criteria, including syntactic dependency and linear order, and the asymmetrical relations are then defined in terms of the Arg1 and Arg2 labels (for example, in Cause:Reason, Arg2 is the cause and Arg1 the effect, while in Cause:Result, Arg1 is the cause and Arg2 is the effect). GraphBank, on the other hand, utilizes a different mechanism to capture the asymmetry. Rather than making reference to linear order, it makes use of directed arcs in the annotation, with definitions provided for how to interpret the directionality for each relation type (for example, for the relation Cause-Effect, the arc is directed from the span stating the cause to the span stating the effect; for the relation Violated Expectation, the arc is directed from the span stating the cause to the span stating the absent effect; and so on).

In the DR-core specification, representation of asymmetry abstracts over the linear ordering and syntactic structure, not only because these are not semantic in nature but also because they may not be good criteria from the viewpoint of interoperability, given the wide variation in cross-linguistic syntax, including clause-combination. Instead, *asymmetry is represented by specifying the argument roles in the definition of each relation*. Arguments are named Arg1 and Arg2, but they bear relation-specific semantic roles. For example, in the Cause relation, defined as “Arg2 serves as an explanation for Arg1” (see [Table 1](#)), the text span named Arg2 always provides the reason in the Cause relation, irrespective of linear order or syntax, and Arg1 always constitutes the result. For human annotators, mnemonic labels indicating the semantic roles, like “reason” and “result”, are more convenient than “Arg2” and “Arg1”, therefore the ISO specification also allows the use of these semantic role labels. [Table 2](#) provides the mapping between Arg1 and Arg2 labels and the corresponding semantic role labels for asymmetric relations. *In symmetric relations, on the other hand, where both arguments play the same semantic role, arguments are named Arg1 and Arg2 following their linear order in the text.*

It is important to note that this representation can be effectively mapped to other schemes for representing asymmetry and in no way obfuscates the differences in linear ordering of the arguments, which can be easily determined by pairing the argument roles with the text span annotations. The ISO scheme acknowledges that linear ordering has a bearing for claims that different versions of an asymmetric relation may not have the same linguistic constraints, for example, with respect to linguistic predictions for the following discourse.^[3]

4.8 Representation of the relative importance of arguments for discourse meaning/structure

Beyond the representation of asymmetry, some frameworks, namely RST, HTDC, and SDRT also explicitly represent the “relative importance” of DRel arguments, taking this relative importance to impact the meaning or structure of the text as a whole. In RST, one argument of an asymmetric relation is labelled the “nucleus” whereas the other is labelled “satellite”, based on the following criteria:^[40] (a) The nucleus is more essential to the writer’s purpose than the satellite; (b) In comparison to the nucleus, the satellite is more easily substitutable without much change to the apparent function of the text (or discourse) as a whole, and (c) Without the nucleus, the content of the satellite is incomprehensible (in the text as a whole), a non sequitur. HTDC has a similar approach, using the term “dominance”, with the goal of deriving a single assertion from a discourse relation connecting two segments, and distinguishing relations in terms of how this single assertion should be derived. In subordinating relations, in particular, the assertion associated with the relation is obtained from the “dominant” segment, as specified in the relation definitions. SDRT, on the other hand, classifies a relation as “subordinating” or “coordinating”,

depending on what structural configuration the arguments create in the discourse graph.^[4] *In the DR-core specification, the relative importance of arguments for the text (meaning or structure) as a whole is not represented directly.* However, because of the explicit identification of the roles of arguments in each relation definition (as described in 4.7), a layer of representation capturing the arguments' relative importance can be easily derived. For example, a mapping from ISO categories to RST categories for Cause would label the Arg2 (corresponding to the reason) argument as the satellite and the Arg1 (corresponding to the result) argument as the nucleus, because there is a one-to-one mapping in RST between the semantic roles of arguments and their respective functional roles for relative importance, for each relation. A similar mapping can be shown for SDRT relations as well.

4.9 Arity of arguments

Except RST, all frameworks assume that a discourse relation has two and only two arguments. In RST, the constraints on the number of arguments for a relation are captured via multinuclear relations; the relations Joint and Sequence (among others) allow for more than two arguments. *In the DR-core specification, a discourse relation is restricted to two and only two arguments,* with the understanding that a mapping from binary relations to n-ary relations is possible where necessary. For example, two identical binary relations with shared arguments, R(A, B) and R(B, C), can be collapsed into a single ternary relation R(A, B, C), if the given framework allows for the relation R to be n-ary.

4.10 Syntactic form, extent, and (non-)adjacency of argument realizations

Three important considerations for annotating the arguments of a discourse relation are the following. The first has to do with the kinds of syntactic forms the realization of an argument can have. That is, what are the minimal allowable syntactic units corresponding to an argument? While all frameworks agree that the *typical* syntactic realization of an argument is a "clause", some allow for certain non-clausal phrases as well. In the end, the differences emerge because of different views on the information status of different syntactic forms in discourse and their relevance to discourse coherence. Also to be considered are languages like Turkish where nominalizations (noun phrases denoting eventualities) are very common.^[79] *In the DR-core specification, what counts as a DRel argument is constrained by its semantic status rather than its syntactic form.* In particular, a DRel argument must denote a situation as defined in 3.2, that is, the situation must be one of the following types: event, state, fact, proposition, or dialogue act^{[2][68]}.

The second issue has to do with the extent of arguments. All frameworks allow for argument spans to be arbitrarily complex, composed of multiple clauses in coordination or subordinate relations, as well as multiple sentences, as long as they are required for interpreting the relation in which they participate. PDTB further stipulates that argument spans must contain the "minimal" amount of information needed to interpret the relation, which is closely related to the third issue concerning the (non-)requirement for the adjacency of argument spans. Some frameworks, such as RST, require the text spans of the related arguments to be textually adjacent, whereas others such as the PDTB impose this constraint only for implicit discourse relations. To a large extent, these differences arise because of differences in assumptions about the global structure of a text, and the reflection of such assumptions in the annotation. As with the issue of syntactic form, it is difficult to reconcile these differences. However, in contrast to the constraint specification for syntactic form, *the DR-core specification remains neutral on the issues of the extent and adjacency of argument spans and does not specify any constraints.* It is important to note, however, that for a fully interoperable annotation scheme, consensus-based constraints must be established for these arguments related features as well. These issues deserve further study and will be addressed in the envisaged second part of the project in which DR-core has been developed.

4.11 Triggers of discourse relations

It is generally agreed that DRels can be realized explicitly in text but can also be implicit, as illustrated in Examples 1 to 4. When explicit, the phrases are typically found to belong to well-defined syntactic classes, such as subordinating conjunctions, coordinating conjunctions, adverbials and prepositional phrases. But some frameworks such as the PDTB also allow for DRels to be realized with other phrase types that don't necessarily correspond to a single syntactic class,^[55] such as the subject-verb sequence

in Example 3. Following this idea, Reference [55] distinguishes DRel expressions based on whether they are frozen closed-class expressions or more productive expressions allowing substitution (cf. “A major reason is” vs. “A most convincing reason is”). Indeed, many connectives from the well-defined and commonly recognized syntactic classes can also be said to be productive if one considers the possibility of their modification (cf. “because” vs. “at least presumably because”). The treatment of modification and negation of discourse relations, as in “not because” and “perhaps because”, is beyond the scope of the DR-core annotation scheme. ISO 24617-6 mentions the possibility of applying qualifiers like “uncertain”, introduced for the annotation of dialogue acts in ISO 24617-2, also to discourse relations and semantic roles. This issue is expected to be taken up in the follow-up of the DR-core project (see 4.16).

Aside from the question of which expression types are taken to be DRel triggers, frameworks still differ in whether the annotation scheme includes the explicit identification or marking of DRel triggers or not. In this respect, only the PDTB and RST currently include the marking of the explicit triggers of discourse relations. In the PDTB, the triggers are marked during the annotation of the discourse relations, whereas in RST, the triggers are added as an additional annotation layer after the annotation of the discourse relations.[70]

With respect to implicit DRels, frameworks differ in whether these inferences are allowed only in adjacent contexts or also in non-adjacent contexts. Here, the framework of GraphBank stands out as the only one to allow implicit DRels between non-adjacent units.

In the DR-core specification, it is considered important to explicitly mark expressions seen as the textual triggers of DRels, since these are valuable clues for inducing models for discourse processing. However, the scheme is flexible about the inference sites for implicit DRels, that is, whether implicit DRels are allowed only between adjacent discourse units or between non-adjacent units as well.

The representation of implicit DRels can also include the insertion of a connective that *could* have been used to express the inferred relation. While all frameworks agree that this is possible, only the PDTB explicitly includes such insertions in its annotation scheme. *In the DR-core specification, insertion of connectives to express inferred DRels is allowed, but not required.*

4.12 Representation of attribution as a discourse relation

Attribution is a relation between agents and situations[77][57] and in many text genres, especially newswire, is observed to occur frequently and in close syntactic interaction with discourse relations and their arguments. In some cases, the relation and its arguments may be attributed to the writer (Example 12) or some other agent introduced in the text (Example 13); in other cases, the relation is established by the writer, with one or both arguments attributed to others (Examples 14 and 15).

NOTE Explicit attributions in the text are shown in courier font for illustration.

Example 12 Since the British auto maker became a takeover target last month, its ADRs have jumped about 78 %.

Example 13 “The public is buying the market when in reality there is plenty of grain to be shipped,” said Bill Biedermann, Allendale Inc. director.

Example 14 Factory orders and construction outlays were largely flat in December while purchasing agents said manufacturing shrank further in October.

Example 15 When Mr. Green won a \$240,000 verdict in a land condemnation case against the State in June 1983, he says Judge O’Kicki unexpectedly awarded him an additional \$100,000.

Although Examples 13 to 15 suggest that the attributions in these sentences do not contribute to the discourse relations identified therein, the close textual coupling of the two has led almost all annotation schemes to annotate attributions in one way or another. Perhaps motivated by the need to not leave any part of the text unconnected, frameworks such as RST, SDRT and GraphBank have in fact treated attribution as a discourse relation, marking a relation called “Attribution” between the attribution phrase (including the agent and attributive predicate) and the attributed content. It is worth noting that none of the original discourse coherence theories on which these annotation frameworks

are based suggest attribution as a discourse relation. Indeed, the original RST formulation argues against the treatment of attributions as rhetorical relations: “Passages that present who said what or attribute information to certain sources rarely relate to other text spans in such a way that relational propositions arise”^[40]. Attribution is annotated in the PDTB as well, but it is not regarded as a discourse relation. Rather, the goal of annotating attribution in the PDTB is to capture semantic interactions of attributions with discourse relations,^[57] the most striking of which is exemplified by Example 16. In this example, the negation associated with the attribution phrase, “I don’t think”, is interpreted lower, with Arg2 of the Contrast relation, so that Arg2 should then be read as “it’s not a main consideration”. Thus, while treating attribution as a different kind of relation, its annotation in the PDTB scheme allows it to be factored out of the discourse relation description while at the same time utilizing its semantic interactions with DRels, via the use of features as described in Reference ^[57].

Example 16 “Having the dividend increases is a supportive element in the market outlook, but I don’t think **it’s a main consideration**,” he says.

The DR-core specification does not treat attribution as a discourse relation, but also does not provide for its annotation. According to ISO 24617-6, the annotation of attribution is recommended as a separate layer, to be undertaken according to a separate annotation scheme. Interactions of attributions with discourse relations can then be utilized by merging the two layers of annotation. Work in this direction can draw on schemes developed in the context of discourse relation annotation,^[57] as well as schemes focusing exclusively on attribution.^{[51][52][53][77]}

4.13 Representation of entity-based relations

There are a few types of connections between segments in discourse that don’t involve a relation between situations in the same way as for other relations such as Cause or Concession. These are connections that different frameworks have variously called “Entity-Elaboration”, “Object-Attribute Elaboration”, “Continuation”, “Circumstance”, “Background”, “Ground-Figure”, “Frame”, “EntRel”, etc. Although there are some fine-grained differences in the exact semantic description across frameworks, they all refer to how the relation affects the narrative or the flow of discourse rather than to a direct relation between the situations denoted by the segments. Moreover, recognition of the connection seems to rest on recognizing a coreferential link between the segments, with one segment providing a description or attribute about an entity mentioned in the other. A striking feature of such relations, that also sets them apart from other relations, is that in all languages studied so far, they defy expression with any kind of connective. Examples 17 to 20 illustrate such relations. In each example, the coreferential entities are highlighted in boldface, and the rough inference of the relation in each case is that the second sentence says something *more* about the co-referring entity. No stronger relation is inferred between the segments, and no connective can be inserted between them.

Example 17 Traders said Goldman Sachs, **Shearson Lehman Hutton** and Salomon Brothers were the main force behind the futures buying at the pivotal moment.

Shearson Lehman Hutton declined to comment.

Example 18 Among the new issues was Massachusetts’s \$230 million of **general obligation bonds**.

The bonds were won by a Goldman, Sachs and Co. group with a true interest cost of 7,17 %.

Example 19 Shortly after the UAL opening, program traders started selling stocks in the **Major Market Index** and S&P 500 index.

The 20-stock **MMI** mimics the Dow Jones Industrial Average.

Example 20 Adding to the junk market’s jitters were reports that Donaldson, Lufkin and Jenrette Securities Corp. is having trouble structuring a \$1,6 billion offering for **TW Food Services Inc.** and will postpone or even cancel the issue.

TW is the largest franchisee of Hardee’s, a fast-food restaurant, and operates several other food chains.

Discussing such relations in the RST Treebank, in particular the “Object-Attribute” relation, it has been argued^[35] that they are orthogonal to the level at which discourse relation structures (i.e. RST trees) are described. The PDTB, following Reference [35], takes the same approach, classifying all such relations as EntRel (short for “Entity-based coherence relations”) and, notably, does not include EntRel in its set of discourse relations. Other frameworks, on the other hand, such as RST, SDRT, and HTDC, have opted to classify such relations as discourse relations. *In the DR-core specification, the latter position is adopted, with the view that these relations, despite being somewhat “functional” in nature, are integral to an understanding of the complete narrative of the discourse.* Furthermore, they are also very frequent in some text genres: in the PDTB, for example, EntRel relations constitute 12 % of the annotated relations. However, exactly how these entity relations contribute to the narrative of the discourse and what the most accurate semantic description would be is something on which a consensus hasn’t been fully established. Consequently, *the current 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.

4.14 Representation of non-existence of a discourse relation

In talking of relations between discourse segments, it is common to assume that in a coherent discourse, every segment of the discourse is related to some other segment. Some annotation frameworks, however, especially those such as the PDTB, where sentential adjacency (as a trigger for implicit relations) and minimality (for selection of argument spans) are important features of the guidelines, the possibility of adjacent segments not being related in any way does arise, though rarely. *In the DR-core specification, a relation is allowed to be marked as “NoRelation” under these conditions.*

4.15 Summary: Assumptions of the DR-core annotation scheme

In summary, the following provides the basic concepts and assumptions underlying the specification of the DR-core annotation scheme for discourse relations.

- a) A discourse relation is a relation expressed in text/speech between situations as abstract semantic objects, such as events, states, facts, propositions, and dialogue acts.
- b) The meaning of discourse relations is described primarily in “informational” terms. Pragmatic aspects of meaning involving beliefs or dialogue acts as one or both of the arguments are represented as a property of arguments, rather than relations.
- c) Discourse relations can be realized explicitly or they can be implicit.
- d) Arguments of discourse relations need not be fully explicit. As per b), beliefs and dialogue acts, when they constitute an argument, may have to be inferred from the realizations.
- d) Discourse relations are categorized as a “flat” set of relations.
- e) Annotations are at a “low level”; the DR-core scheme only describes the nature of the low-level structure of a discourse.
- f) Explicit triggers of discourse relations can be represented as a feature of the relation.
- g) A discourse relation takes two and only two arguments, labelled as “Arg1” and “Arg2”. An argument may be realized by multiple clauses or other grammatical units.
- h) Asymmetrical relations are represented with relation-specific argument role labels that are tied to the definitions provided for the relation. For each asymmetric discourse relation the DR-core specification also provides mnemonic names of its arguments instead of “Arg1” and “Arg2”; these mnemonic names may alternatively be used in annotations. For a symmetrical relation, the labels “Arg1” and “Arg2” follow their order of occurrence in the discourse.
- i) Two given situations can be related by more than one relation.

- j) The relative importance of a relation's arguments with respect to the text as a whole (cf. subordinating and coordinating relations in SDRT, nucleus and satellite in RST) is not represented as such. No a priori assumptions are made concerning constraints on syntactic form, syntactic complexity, or textual adjacency of expressions that may correspond to the arguments of a discourse relation.
- k) Attribution relations are not treated as discourse relations and are assumed to be provided as a separate layer of annotation.
- l) The annotation of implicit relations may optionally include the specification of a connective that could express the inferred relation. To accommodate annotation frameworks that impose strict adjacency constraints on the marking of relations, the label "NoRelation" can be utilized when two adjacent segments are not related by any discourse relation.

These choices concerning the basic concepts involved in the annotation of discourse relations and the way these concepts are related are reflected in the metamodel shown in [Figure 1](#).

4.16 Issues to be taken up in the follow-up of DR-core

A future part of ISO 24617 will address several key issues, including but not limited to, the following:

- a) constraints on extent and adjacency requirements for argument spans (see [4.10](#));
- b) constraints on the inference sites for implicit discourse relations (see [4.11](#));
- c) hierarchical classification of discourse relations (see [4.5](#));
- d) constraints on types of expressions that count as DRel triggers (see [4.11](#));
- e) clarification of situation types and their definitions (see [Clause 5](#));
- f) clarification of the Expansion relation and its possible refinements (see [4.13](#));
- g) treatment of negated and modified discourse relation (see [4.12](#));
- h) assessment of the applicability of the standard to multiple languages, and extensions to the standard as necessitated by this study.

4.17 Metamodel

Discourse relations and their arguments take central stage in the metamodel of the DR-core. The metamodel reflects the assumptions a) and b) listed in [4.15](#) by the fact that relation arguments can have different types, as indicated by the link from relation arguments to argument types. It is assumed that the set of argument types includes the types of the situations that make up sentence meanings [including events, states, processes, facts, conditions, as well as negated eventualities (as in "Mary smiled at John, *but he didn't smile back*"), for which the generic term "situations" is often used], as well as beliefs and dialogue acts involved in utterance meanings and in "pragmatic" interpretations of discourse relations (as in "Carl is a fool; he beats his wife").

The metamodel incorporates assumption d), that the arguments of a discourse relation are always realized explicitly by the fact that each argument is related to one markable (represented by the number "1" at the tip of the arrow from arguments to markables; see [Figure 1](#)), which in turn is associated with one segment of primary data. The fact that a discourse relation can be explicit or implicit [assumption c)] is reflected in the indication "0...1" at the tip of the arrow from discourse relations to markables.




Figure 1 — Metamodel for the annotation of semantic relations in discourse

The assumptions h) and i) are represented in the metamodel by the number “2” at the tip of the arrow from discourse relations to arguments, and also by the “2” at the tip of the arrow connecting discourse relations with their argument roles. The dotted arrows at the bottom of [Figure 1](#) indicate possible links to another layer of annotation, concerned with the identification of the source to which a discourse relation or (one or both of) its arguments may be attributed [assumption l)].

5 Core discourse relations

[Table 1](#) lists the set of core discourse relations of the DR-core annotation scheme. The level of granularity is motivated by the consideration that the chosen relations encapsulate the minimal description that has been more or less successfully implemented in various annotation efforts to date. This set is by no means fixed and can be augmented in three ways: (1) with more fine-grained refinements of the core relations; (2) with less fine-grained abstractions of the core relations; and (3) with equally fine-grained relations that may supplement the set of core relations.

Table 1 — Relations and argument roles

	ISO 24617-8	Symmetry	Relations and argument role definitions
1.	Cause	Asymmetric	Arg2 is an explanation for Arg1.
2.	Condition	Asymmetric	Arg2 is an unrealized situation which, when realized, would lead to Arg1.
3.	Negative condition	Asymmetric	Arg2 is an unrealized situation which, when “not” realized, would lead to Arg1.
4.	Purpose	Asymmetric	Arg2 is the goal or purpose of the situation described by Arg1.
5.	Manner	Asymmetric	Arg2 describes how Arg1 comes about or occurs.

Table 1 (continued)

	ISO 24617-8	Symmetry	Relations and argument role definitions
6.	Concession	Asymmetric	An expected causal relation between Arg1 and $\neg$ Arg2 is cancelled or denied by Arg2.
7.	Contrast	Symmetric	One or more differences between Arg1 and Arg2 are highlighted with respect to what each predicates as a whole or to some entities they mention.
8.	Exception	Asymmetric	Arg2 indicates one or more circumstances in which the situation(s) described by Arg1 does not hold.
9.	Similarity	Symmetric	One or more similarities between Arg1 and Arg2 are highlighted with respect to what each predicates as a whole or to some entities they mention.
10.	Substitution	Asymmetric	Arg1 and Arg2 are alternatives, with Arg2 being the favoured or chosen alternative.
11.	Conjunction	Symmetric	Arg1 and Arg2 bear the same relation to some other situation evoked in discourse. Their conjunction indicates that they both hold with respect to that situation.
12.	Disjunction	Symmetric	Arg1 and Arg2 bear the same relation to some other situation evoked in the discourse, explicitly or implicitly. Their disjunction indicates that they are alternatives with respect to that situation, with the disjunction being non-exclusive so that both Arg1 and Arg2 may hold.
13.	Exemplification	Asymmetric	Arg1 describes a set of situations; Arg2 describes an element of that set.
14.	Elaboration	Asymmetric	Arg1 and Arg2 are the same situation, but Arg2 contains more detail.
15.	Restatement	Symmetric	Arg1 and Arg2 describe the same situation, but from different perspectives.
16.	Synchrony	Symmetric	Some degree of temporal overlap exists between Arg1 and Arg2. All forms of overlap are included.
17.	Asynchrony	Asymmetric	Arg1 temporally precedes Arg2.
18.	Expansion	Asymmetric	Arg2 is a situation involving some entity/entities in Arg1, expanding the narrative forward of which Arg1 is a part, or expanding on the setting relevant for interpreting Arg1. The Arg1 and Arg2 situations are distinct.
19.	Functional dependence	Asymmetric	Arg2 is a dialogue act with a responsive communicative function; Arg1 is the dialogue act(s) that Arg2 responds to.
20.	Feedback dependence	Asymmetric	Arg2 is a feedback act that provides or elicits information about the understanding or evaluation by one of the dialogue participants of Arg1.

The DR-core set of core discourse relations is based on the establishment of semantic equivalences with five well-known semantic taxonomies for discourse relations: PDTB[46]; Kehler's classification of coherence relations,[32][33][34] largely based on HTDC 9,[17][18] RST,[40] SDRT[3] and CCR.[66] It also draws on the experiences with discourse relation annotation in multiple languages and genres,[10][78][58][49][60][80][82][48][1][65] among others).

Table 2 — Mnemonic names for argument roles of asymmetric discourse relations

	Discourse relation	Arg1	Arg2
1.	Cause	Result	Reason
2.	Condition	Consequent	Antecedent
3.	Negative Condition	Consequent	Negated-antecedent
4.	Purpose	Enablement	Goal

Table 2 (continued)

	Discourse relation	Arg1	Arg2
5.	Manner	Achievement	Means
6.	Concession	Expectation-raiser	Expectation-denier
7.	Exception	Regular	Exclusion
8.	Substitution	Disfavoured-alternative	Favoured-alternative
9.	Exemplification	Set	Instance
10.	Elaboration	Broad	Specific
11.	Asynchrony	Before	After
12.	Expansion	Narrative	Expander
13.	Functional dependence	Antecedent-act	Dependent-act
14.	Feedback dependence	Feedback-scope	Feedback-act

As discussed in [Clause 6](#), because of different linear ordering possibilities of argument realizations, annotating the Arg1 and Arg2 arguments of asymmetric relations relies heavily on knowing the semantic role of each of the arguments, as provided in the definition of each relation (see [Table 1](#), column 4). While some frameworks have represented asymmetry using syntactic criteria, the DR-core specification represents this semantically, by having the argument labels Arg1 and Arg2 bear semantic roles defined for each relation. Mnemonic names of these semantic roles are listed in [Table 2](#) and may be used as alternatives to “Arg1” and “Arg2”, to which they can be mapped using this table.

Definitions and examples for each relation are provided below. In all definitions, the arguments of the relation, Arg1 and Arg2, refer to situations, which are abstract semantic objects of the following types: event, state, fact, proposition, belief and dialogue act. In all examples, explicit or implicit (inserted) expressions conveying the relation are underlined, with inserted expressions further enclosed in parentheses. Argument realizations are shown with *italic face* for Arg1, and **bold face** for Arg2. It must be noted that the argument spans in these examples follow the guidelines of the PDTB[61] but this is just one possible way to mark argument spans. As noted in [4.10](#), the DR-core scheme is underspecified with respect to constraints on the selection of argument realization spans. Most examples are taken from Reference [43]. Where possible, examples are provided to illustrate both relations realized explicitly and implicit relations, as well as relations involving dialogue acts and beliefs, which are marked with the prefixes [+DA] and [+PA] for the relevant arguments.

— **Cause:**

Def: Arg2 (Reason) is an explanation for Arg1 (Result).

Symmetry: Asymmetric

Examples:

- Perhaps because **they won**, Mr. Bork's attackers come through more vividly than his defenders.
- Sears is negotiating to refinance its Sears Tower for close to \$850 million, sources said. (Implicit = because) **The retailer was unable to find a buyer for the building.**
- Now, though, enormous costs for earthquake relief will pile on top of outstanding costs for hurricane relief.** *“That obviously means that we won't have enough for all of the emergencies that are now facing us, and we will have to consider appropriate requests for follow-on funding,”* Mr. Fitzwater said.
- [+PA] *The nations of southern Africa know a lot about managing elephants;* (Implicit = as) **their herds are thriving.**
- [+DA] *What makes people blurt out their credit-card numbers to a caller they've never heard of? Do they really believe that the number is just for verification and is simply a formality on the road to being a grand-prize winner? What makes a person buy an oil well from some stranger knocking*

on the screen door? Or an interest in a retirement community in Nevada that will knock your socks off, once it is built? Because in the end, these people always wind up asking themselves the same question: "How could I be so stupid?"

f) [+DA] *My favourite is the police department. (Implicit = because) **They're not aimed at the criminal.*** The judicial system is aimed at the citizens.

g) U: I thought well from like Elmira through to Corning to Bath is how many hours?

S: four hours

U: so [+DA] *we're like screwed as far as those two box cars at Bath.*

— **Condition:**

Def: Arg2 (Antecedent) is an unrealized situation which, when realized, would lead to Arg1 (Consequent).

Symmetry: Asymmetric

Examples:

- a) But some bond market analysts said *that could quickly change* if property casualty insurance companies scramble to sell portions of their municipal portfolios to raise cash to pay damage claims.
- b) If anyone has difficulty imagining a world in which history went merrily on without us, [+DA] *Mr. Gould sketches several.*

— **Negative Condition:**

Def: Arg2 (Negated Antecedent) is an unrealized situation which, when *not* realized, would lead to Arg1 (Consequent).

Symmetry: Asymmetric

Examples:

- a) Unless the Federal Reserve eases interest rates soon to stimulate the economy, *profits could remain disappointing.*
- b) *But a Soviet bank here would be crippled* unless Moscow found a way to settle the \$188 million debt, which was lent to the country's short-lived democratic Kerensky government before the Communists seized power in 1917.
- c) If you cannot give me a fully convincing explanation here and now, [+DA] *you're fired.*

— **Purpose:**

Def: Arg2 (Goal) is the goal or purpose of the situation described by Arg1 (Enablement).

Symmetry: Asymmetric

Examples:

- a) *Skilled ringers use their wrists to advance or retard the next swing, so that one bell can swap places with another in the following change.*
- b) Adjusters must count the number of bathrooms, balconies, fireplaces, chimneys, microwaves and dishwashers. **But they must also assign a price to each of these items as well as to floors, wall coverings, roofing and siding, to come up with a total value for a house.** *To*

do that, they must think in terms of sheetrock by the square foot, carpeting by the square yard, wallpaper by the roll, moulding by the linear foot.

— **Manner:**

Def: Arg2 (Means) specifies how Arg1 (Achievement) comes about or occurs.

Symmetry: Asymmetric

Examples:

- a) US trade negotiators argue *that countries with inadequate protections for intellectual-property rights could be hurting themselves* by **discouraging their own scientists and authors...**
- b) *The secretary spoke as if* **he were the CEO.**

— **Concession:**

Def: An expected causal relation between Arg1 (Expectation-raiser) and $\neg$ Arg2 (Expectation-denier) is cancelled or denied by Arg2.

Symmetry: Asymmetric

Examples:

- a) *Mr. Coleman said this week that he would devote the remainder of the political season to positive campaigning, but* **the truce lasted only hours.**
- b) [+DA] *Here in the Atlanta area, our crime rate is just astronomical, yet, you go on the streets and they're giving speeding tickets.*

— **Contrast:**

Def: One or more differences between Arg1 and Arg2 are highlighted with respect to what each predicates as a whole or to some entities they mention.

Symmetry: Symmetric

Examples:

- a) *The prices of puts generally didn't soar Friday. For example, the premium as a percentage of the stock price for certain puts on Eli Lilly and Co. moved up from 3 % at Thursday's close to only 3.3 % at Friday's close, even though the shares dropped more than \$5.50. But* **put-option prices may zoom when trading resumes today.**
- b) Now, neither side – the US nor Mr. Noriega – has an easy out. *President Bush has sworn to bring him to justice. (Implicit = on the other hand)* **Mr. Noriega believes he hasn't any alternative but to continue clutching to power.** It is a knockout battle, perhaps to the death.

— **Exception:**

Def: Arg2 (Exclusion) evokes one or more circumstances in which Arg1 (Regular) does not hold.

Symmetry: Asymmetric

Examples:

- a) *Boston Co. officials declined to comment on Moody's action on the unit's financial performance this year except to deny a published report that outside accountants had discovered evidence of significant accounting errors in the first three quarters"* **results.**

- b) Michael Ross, a New York lawyer who heads the ABA's grand jury committee, said *that lawyers are prohibited by the ABA's code of ethics from disclosing information about a client except where a court orders it or to prevent the client from committing a criminal act that could result in death.*

— **Similarity:**

Def: One or more similarities between Arg1 and Arg2 are highlighted with respect to what each predicates as a whole or to some entities they mention.

Symmetry: Symmetric

Examples:

- a) *Just as the 1980s bull market transformed the US securities business, so too will the more difficult environment of the 1990s,*" says Christopher T. Mahoney, a Moody's vice president.
- b) *Compromises are possible. Citizens in Peninsula, Ohio, upset over changes to a bridge, negotiated a deal: The bottom half of the railing will be type F, while the top half will have the old bridge's floral pattern. Similarly, highway engineers agreed to keep the old railings on the Key Bridge in Washington, DC, as long as they could install a crash barrier between the sidewalk and the road.*

— **Substitution:**

Def: Arg1 (Disfavoured-alternative) and Arg2 (Favoured-alternative) are alternatives, with Arg2 being the favoured or chosen alternative.

Symmetry: Asymmetric

Examples:

- a) *Rather than increasing dividends, some companies have used cash to buy back some of their shares,* notes Steven G. Einhorn, co-chairman of the investment policy committee at Goldman, Sachs and Co.
- b) *Ground zero of the HUD scandal is the Secretary's "discretionary fund," a honey pot used to fund projects that weren't approved through normal HUD channels. Jack Kemp wants to abolish it. Instead, Congress's idea of reform is to increase this slush fund by \$28.4 million.*
- c) [+DA] *My favourite is the police department. They're not aimed at the criminal. The judicial system is aimed at the citizens.*

— **Conjunction:**

Def: Arg1 and Arg2 bear the same relation to some situation evoked in the discourse, explicitly or implicitly. Their conjunction indicates that they both hold with respect to that situation.

Symmetry: Symmetric

Examples:

- a) *But investors trying to play all the angles may find that stock splits are a lot like cotton candy: They look tempting, but there's hardly any substance. And they can even leave a sticky problem, in the form of higher brokerage commissions.*
- b) *"When we evaluated raising our bid, the risks seemed substantial and persistent over the next five years, and the rewards seemed a long way out. That got hard to take," he added.*

— **Disjunction:**

Def: Arg1 and Arg2 bear the same relation to some “situation evoked in the discourse, explicitly or implicitly. Their disjunction indicates that they are alternatives with respect to that situation. The disjunction is non-exclusive so Arg1 and Arg2 may both hold.

Symmetry: Symmetric

Examples:

- a) “You’ve either got a chair or you don’t.”
- b) If we want to support students, we might adopt the idea used in other countries of offering more scholarships based on something called “scholarship,” rather than on the government’s idea of “service.” Or we might provide a tax credit for working students.
- c) Under two new features, participants will be able to transfer money from the new funds to other investment funds or, if their jobs are terminated, receive cash from the funds.

— **Exemplification:**

Def: Arg1 (Set) is a set of situations; Arg2 (Instance) is an element of that set.

Symmetry: Asymmetric

Examples:

- a) The computers were crude by today’s standards. **Apple II owners, for example, had to use their television sets as screens and stored data on audiocassettes.**
- b) Publishing officials believe that while Random House has enjoyed spectacular growth and has smoothly integrated many acquisitions in recent years, some of the bigger ones haven’t been absorbed so easily. (Implicit = for example) Crown Publishing Group, acquired last year, is said **to be turning in disappointing results.**
- c) Typically, these laws seek to prevent executive branch officials from inquiring into whether certain federal programs make any economic sense or proposing more market-oriented alternatives to regulations. **Probably the most egregious example is a proviso in the appropriations bill for the executive office that prevents the president’s Office of Management and Budget from subjecting agricultural marketing orders to any cost-benefit scrutiny.**

— **Elaboration:**

Def: Arg1 (Broad) and Arg2 (Specific) are the same situation, but Arg2 provides more detail.

Symmetry: Asymmetric

Examples:

- a) **Many modern scriptwriters seem to be incapable of writing drama, or anything else, without foul-mouthed cursing. Sex and violence are routinely included even when they are irrelevant to the script, and high-tech special effects are continually substituted for good plot and character development.** *In short, we have a movie and television industry that is either incapable or petrified of making a movie unless it carries a PG-13 or R rating.*
- b) *An enormous turtle has succeeded where the government has failed:* **(Implicit = Specifically) He has made speaking Filipino respectable.**

- c) [+DA] *The judges sitting in that kind of stuff day after day, they know all the procedures, they know what's good and what isn't, **they'd be able to say, I'm sorry, you can't use this as evidence.***

— **Restatement:**

Def: Arg1 and Arg2 are the same situation, but viewed from different perspectives.

Symmetry: Symmetric

Examples:

- a) *Some days the coaches make you feel as though you are part of a large herd of animals. In other words, **they treat you like a piece of meat.***
- b) *Yet some people are advancing a chilling casuistry: *that what we are seeing is somehow the understandable result of the historical sins committed by the Turks in the 16th century.**

Today's Turks in Bulgaria, in other words, deserve what is coming to them four centuries later.

— **Synchrony:**

Def: Some degree of temporal overlap exists between Arg1 (Before) and Arg2 (After).

All forms of overlap are included.

Symmetry: Symmetric

Examples:

- a) *The company is operating under Chapter 11 of the federal Bankruptcy Code, *giving it court protection from creditors' lawsuits* while it attempts to work out a plan to pay its debts.*
- b) *Then, in late-afternoon trading, hundred-thousand-share buy orders for UAL hit the market, including a 200,000-share order through Bear Stearns that seemed to spark UAL's late price surge. Almost simultaneously, PaineWebber began a very visible buy program for dozens of stocks.*
- c) *The parishioners of St. Michael and All Angels stop to chat at the church door, as members here always have. (Implicit = while) **In the tower, five men and women pull rhythmically on ropes attached to the same five bells that first sounded here in 1614.***

— **Asynchrony:**

Def: Arg1 temporally precedes Arg2.

Symmetry: Asymmetric

Examples:

- a) *A buffet breakfast was held in the museum, where food and drinks are banned to everyday visitors. Then, in the guests' honour, **the speedway hauled out four drivers, crews and even the official Indianapolis 500 announcer for a 10-lap exhibition race.***
- b) *William Gates and Paul Allen in 1975 developed an early language-housekeeper system for PCs, **and Gates became an industry billionaire** six years after IBM adapted one of these versions in 1981.*
- c) *The Artist has his routine. *He spends his days sketching passers-by, or trying to.* (Implicit = then) **At night he returns to the condemned building he calls home.***
- d) [+DA] ***So, we first go down?** After passing the wild heath?*

— **Expansion:**

Def: Arg2 (Expander) is a situation involving some entity/entities in Arg1 (Narrative), expanding the narrative of which Arg1 is a part, or expanding on the setting relevant for interpreting Arg1. The Arg1 and Arg2 situations are distinct.

Symmetry: Asymmetric

Notes: Typically, these relations are implicit and do not allow insertion of a connective to express the relation. In English, “and” can be used when Expansion carries the narrative forward.

Examples:

- a) *Anticipating the Fed's move, money traders lowered a key interest rate known as the Federal Funds rate to 8.625 % late Friday, down from 8.820 % the day before. (Implicit-Expansion) Tiny movements in the rate, which is what banks charge each other for overnight loans, are usually among the few visible tracks that the Fed leaves on the monetary markets.*
- b) *Tandem said it expects to report revenue of about \$450 million and earnings of 35 cents to 40 cents a share. (Implicit-Expansion) The results, which are in line with analysts' estimates, reflect “a continued improvement in our U.S. business,” said James Treybig, Tandem's chief executive officer.*
- c) *Analysts' third-quarter estimates for the Midland, Mich., company are between \$3.20 a share and \$3.30 a share, compared with \$3.36 a year ago, when profit was \$632 million on sales of \$4.15 billion. (Implicit-Expansion) A Dow spokeswoman declined to comment on the estimates.*
- d) [+DA] *So we first go down? After passing the wild heath?*
- e) [+DA] *A: now you continue to the right.*
B: To the right again.
A: Until about half way between the adobe huts and the right border of the map.

— **Functional dependence:**

Def: Arg2 (Dependent-act) is a dialogue act with a responsive communicative function;
 Arg1 (Antecedent-act) is the dialogue act that Arg2 responds to.

Symmetry: Asymmetric

Notes:

- a) Due to its responsive character, the determination of the semantic content of Arg2 in general requires determination of the semantic content of Arg1.
- b) Responsive dialogue acts, defined in ISO 24617-2, are: Agreement, Disagreement, Answer, Confirm, Disconfirm, Correction, Accept Request, Decline Request, Address Request, Accept Offer, Decline Offer, Address Offer, Accept Suggest, Decline Suggest, Address Suggest, Return Greeting, Return Self-Introduction, Accept Apology, Accept Thanking, and Return Goodbye.
- c) This relation cannot be expressed by a discourse connective.

Examples:

- a) *A: What newspapers do you read? [Question]*

B: Well, I read uh the local newspaper, and I also try and read one of the uh major dailies like the Chicago Tribune, or the New York Times or something like that. [Answer]

b) *B: I really like NPR a lot* [Inform]

A: Yeah that's pretty good. [Agreement]

— **Feedback dependence:**

Def: Arg2 is a dialogue act with a feedback function that provides or elicits information about the understanding or evaluation by one of the dialogue participants of Arg1,

Symmetry: Asymmetric

Notes:

- a) The semantic content of Arg2 is mostly provided by Arg1.
- b) The ISO standard for dialogue act annotation (ISO 24617-2) defines five feedback functions: Auto-Positive and Auto-Negative, providing positive and negative feedback information, respectively, about the speaker's processing of something that was said by another speaker; Allo-Positive and Allo-Negative, similarly about the addressee's processing of something that the speaker has said; and Feedback Elicitation, eliciting feedback about the addressee's processing of something that the speaker has said.
- c) This relation cannot be expressed by a discourse connective, but is often expressed in English by "OK", "right", "pardon?" and in similar ways in other languages.

Examples:

- a) **A: What newspapers do you read?**

B: Well, I read uh the local newspaper, and I also try and read one of the uh... major dailies like the Chicago Tribune, or the New York Times or something like that.

A: Uh-huh.

- b) **A: We are going to go due south straight south and then we're going to g— turn straight back round and head north past an old mill on the right-hand side.**

B: Due south and then back up again.

6 Current approaches and annotation schemes

6.1 Overview

This clause provides a discussion of the major existing theories and annotation frameworks for representing discourse relations. The discussion focuses on aspects relevant to the DR-core specification. The clause ends with a mapping provided from the DR-core discourse relations to the relations found within each of the frameworks under discussion. The goal here is to provide a basis for comparison and mappings between frameworks, methods and corpora, using ISO 24617-8 core relations as a pivot representation.

6.2 Rhetorical structure theory (RST)

The goal of RST^[40] is to provide an account of text organization, addressing what counts as the units of a text, what relations hold between the units, and how the units link with each other to form a connected whole. Based on a study of a large number of texts from various genres, and originally conceived for the purpose of building computer systems with text generation capabilities, RST describes relations between textual units (minimally and typically, clauses) in functional terms, that is, in terms of the

writer's assumptions about the reader and the writer's intentions for effects on the reader. It also developed the idea of a hierarchical tree structure representation for texts, with five types of schema (two mononuclear schemas, and three multinuclear) defined for how units can be connected with each other, with adjacency of the units being a strict constraint. Structure building is done in a recursive fashion, with basic units assumed to be primarily clauses. In a mononuclear schema, one argument of the relation is highlighted as the "nucleus", namely the argument that is more essential to the writer's purpose, whereas the other argument(s), called "satellites", are non-essential and would be non sequiturs if not connected with the nucleus argument. Relations in RST are proposed as an "open set", with the 25 relations described in the original work shown in [Table 3](#).

Table 3 — Classical RST relations

RST Subject-matter relations	RST Presentational relations
1. Elaboration	19. Motivation (increases desire)
2. Circumstance	20. Antithesis (increases positive regard)
3. Solutionhood	21. Background (increases ability)
4. Volitional Cause	22. Enablement (increases ability)
5. Volitional Result	23. Evidence (increases belief)
6. Non-volitional Cause	24. Justify (increases acceptance)
7. Non-volitional Result	25. Concession (increases positive regard)
8. Purpose	
9. Condition	
10. Otherwise	
11. Interpretation	
12. Evaluation	
13. Restatement	
14. Summary	
15. Sequence	
16. Contrast	
17. Joint	
18. List	

Reference [\[40\]](#) suggests that relations can be taxonomized in different ways, depending on one's goals, and themselves propose a bipartite taxonomy on the basis of the type of intended effect of the relation on the reader, grouping relations as either "subject-matter" or "presentational", as shown in [Table 3](#). In subject-matter relations, the intended effect is that the reader recognizes the relation in question, whereas in presentational relations, the intended effect is to increase some inclination in the reader.

Besides the original work on RST by Mann and Thompson,[\[40\]](#) RST is revisited by Taboada and Mann. [\[71\]](#)[\[72\]](#)

6.3 RST Treebank

The first large-scale corpus annotated with discourse relations was the RST Discourse Treebank (RST Treebank).[\[10\]](#) Annotated over 385 texts from the WSJ texts in the Penn Treebank, the annotation framework is primarily modelled after RST, with the following differences. First, the relation set used in the RST corpus, shown in the second column of [Table 4](#), is much larger than the 25 relations proposed in the original theory. The table shows further grouping of the relations (second column), as suggested in Reference [\[10\]](#), although this is not directly represented in the annotation. The last grouping as "structural" does not comprise semantic relations, but rather organizational relations designed as conventions for annotation. As described in the RST corpus tagging manual,[\[9\]](#) the annotation scheme reflects the nuclearity status of the relations as well, and further indicates the semantic role of the

arguments. Thus, for example, the relation “Cause” is a mononuclear relation, whose nucleus conveys the cause of the situation described by the satellite. A significant feature of the RST corpus, not shown here, is that the nuclearity status of a relation can vary. For example, the “Evaluation” relation has two versions, one in which the argument providing the evaluation is the nucleus, and another in which it is the satellite. Furthermore, some relations, for example “Evaluation”, can be both mononuclear as well as multinuclear. The second departure of the corpus from the original theory has to do with what is considered an elementary discourse unit (EDU). While RST suggests the idea of a “clause” as a minimal unit, the RST corpus also allows for some non-clausal units, in particular phrases that begin with strong discourse connectives, such as “because of”, “in spite of”, “according to”, etc. Furthermore, clauses that appear as subjects, objects or complements of a verb are not treated as EDUs, while relative clauses, nominal post-modifiers or clauses that break up other legitimate clauses are treated as embedded EDUs. Identification of boundaries was done using lexical and syntactic cues.

Table 4 — RST Treebank relations

RST Treebank Classes	RST Treebank Relations
1. Attribution	Attribution, Attribution-negative
2. Background	Background, Circumstance
3. Cause	Cause, Result, Consequence
4. Comparison	Comparison, Preference, Analogy, Proportion
5. Condition	Condition, Hypothetical, Contingency, Otherwise
6. Contrast	Contrast, Concession, Antithesis
7. Elaboration	Elaboration-additional, Elaboration-general-specific, Elaboration-part-whole, Elaboration-process-step, Elaboration-object-attribute, Elaboration-set-member, Example, Definition
8. Enablement	Purpose, Enablement
9. Evaluation	Evaluation, Interpretation, Conclusion, Comment
10. Explanation	Evidence, Explanation-argumentative, Reason
11. Joint	List, Disjunction
12. Manner-Means	Manner, Means
13. Topic-Comment	Problem-solution, Question-answer, Statement-response, Topic-comment, Comment-topic, Rhetorical-question
14. Summary	Summary, Restatement
15. Temporal	Temporal-before, Temporal-after, Temporal-same-time, Sequence, Inverted-sequence
16. Topic Change	Topic-shift, Topic-drift
17. Structural	Textual-organization, Same-unit

The steps of annotation in the RST corpus involve, first, a segmentation of the full text into EDUs, and second, building the hierarchical tree discourse structure by recursively linking adjacent EDUs and labelling the relationship with a relation from the relation set ([Table 4](#)) while indicating the nuclearity status.

It is worth noting that as in the RST theory, the RST Treebank admits only one relation to hold between EDUs. Although the possibility of multiple relations is acknowledged,^[9] for example, a causal and temporal relation holding simultaneously between two EDUs, annotators were instructed to select only the most salient one, using a predefined criterion for salience-based ranking of the relations.

Besides the RST Discourse Treebank, there are other RST Treebanks for English,^[13] as well as for several other languages, including German,^[69] Portuguese,^[50] Spanish^[14] and Basque.^[15]

6.4 Hobbs' Theory of Discourse Coherence (HTDC)

Hobbs' account of discourse relations (what he calls "coherence relations") and discourse structure is closely tied to a knowledge-based theory of discourse interpretation. Discourse relations are formalized in terms of the inferences drawn by the readers or listeners that allow them to recognize the relations in question. Higher-level structures composed from individual discourse relations and their arguments are essentially trees, though he argues that not all discourses would permit the composition of a single tree over the entire text. The discourse relations defined by Hobbs, shown in the first column of [Table 5](#), comprise a smaller number of relations when compared with other theories.

Table 5 — Relations in HTDC and GraphBank

HTDC	GraphBank
Occasion	Temporal Sequence
Evaluation	Elaboration
Ground-Figure/Background	
Elaboration	
Cause	Cause-effect
Explanation	
Parallel	Similarity
Contrast	Contrast
Exemplification	Example
Generalization	Generalization
Violated Expectation	Violated Expectation
—	Condition
—	Attribution
—	Same-segment

6.5 GraphBank

The GraphBank corpus^[78] takes as its starting point the relation set proposed in Hobbs' work and extends it (second column of [Table 5](#)) with a few new relations added ("Condition", "Attribution", "Same-segment") and a few distinctions collapsed ("Elaboration" and "Cause-effect"), the latter being due to the decision to implement asymmetrical relations in the annotation with directed edges between arguments rather than with different relations. (Note that "Exemplification/Example" and "Generalization" could have been collapsed in a similar manner but this is not done in GraphBank.) The notion of directionality is related to the notion of nucleus and satellite in RST, and for asymmetrical relations, the directionality is from the satellite node to the nucleus node. For each asymmetrical relation, the satellite and nucleus are explicitly identified in terms of the semantic role of the arguments. ^[81] For example, in a "Cause-Effect" relation, the argument expressing the cause is the satellite whereas the argument expressing the effect is the nucleus. The annotation procedure in GraphBank consists of three steps and is roughly modelled along the lines of the process outlined in ^[18]. The steps involve first, segmentation of the text into discourse segments; second, grouping of topically-related segments; and third, identification of coherence relations between segments and groups of segments. Discourse segmentation assumes a clause-like unit as the basic unit, and segment grouping is done according to several criteria, for example, when segments are attributed to the same source or when they are about the same topic or subtopic. Segment grouping itself creates a partially hierarchical structure for the text. Coherence relation recognition is to a large extent guided by recognizing conjunctions or connectives, provided to annotators, used to relate discourse segments, for both when the conjunctions are explicit in the text as well as when they are absent. In the latter case, recognition involves testing the inferred relation by inserting a corresponding conjunction from the pre-defined list. As in RST, the goal in GraphBank is to connect every discourse segment in the text with some other segment, resulting in a single connected structure for the entire text. However, GraphBank importantly makes no a priori assumptions about the nature of the resulting global structure, but rather lets the structures emerge

from the annotations with no constraints, such as segment adjacency imposed on where the discourse relations should or should not be inferred. A study of the resulting annotations in GraphBank leads to the conclusion that chain graphs rather than trees are descriptively more adequate as a data structure for textual discourse.

6.6 SDRT

In the Segmented Discourse Representation Theory (SDRT),^[37] discourse relations are defined in terms of their contribution to truth conditions and are grouped into subordinating or coordinating relations based on how they contribute to building the structure of the discourse. Relations are also classified as either veridical (i.e. relations that entail both of their arguments) and non-veridical (i.e. relations where only one of the arguments is entailed). Using a clause-based definition of the EDU, including appositive and non-restrictive relative clauses, and permitting connections between non-adjacent units, higher-level structures in SDRT are posited as directed acyclic graphs, allowing for multiple parents, multiple attachments, multiple relations between two segments, as well as crossed dependencies. Discourse graph construction in SDRT is constrained by the right frontier principle according to which each new EDU must attach either to the last discourse unit or to some unit in the path from the last unit to the top node in the graph. The theory also allows the creation of complex units from EDUs, and allows discourse relations to take complex units as arguments. Complex units are built up recursively from EDUs and the relations that connect them. The core set of discourse relations proposed in Reference [37] are shown in the first column of Table 6. Additional versions of the core relations have also been proposed,^[37] based on the level at which they are interpreted (content-level, cognitive-level, metatalk-level) as well as versions intended to handle utterances in dialogue, but these are not shown in the Table, since they are not accounted for in the SDRT-related corpora. Corpora modelled after SDRT include the DISCOR corpus,^{[5][64]} annotated over English texts, and the ANNODIS corpus,^[1] annotated over French texts. The annotation procedures for both corpora are largely the same, involving first, segmentation of the text into EDUs and then constructing a discourse graph by connecting EDUs with discourse relations while also recursively building complex discourse units and incorporating them in the discourse graph as well. The set of relations in the corpora includes some omissions from the set proposed in Reference [37], as well as some additions, as shown in the second and third column of Table 6.

Table 6 — Relations in SDRT, DISCOR and ANNODIS

SDRT	DISCOR	ANNODIS
Alternation	Alternation	Alternation
Background	Background	Background
Consequence	Consequence	Conditional
Continuation	Continuation	Continuation
Elaboration	Elaboration	Elaboration
Explanation	Explanation	Explanation
Narration	Narration	Narration
Result	Result	Result
Contrast	Contrast	Contrast
Parallel	Parallel	Parallel
—	Precondition	Flashback
—	Commentary	Comment
—	Attribution	Attribution
—	Source	—
—	—	Entity-elaboration
—	—	Frame
—	—	Temporal-location
—	—	Goal

6.7 CCR

In Reference [66], which focuses on specifying the semantics of discourse relations and does not address the nature of data structures for discourse, discourse relations can be defined and classified in terms of a few cognitively motivated basic principles. The first, “Basic Operation”, classifies a relation as either “Causal” or “Additive”. The second, “Source of Coherence”, classifies a relation as either “Semantic” or “Pragmatic”. The Third, “Order of Segments”, classifies a relation as either “Basic” or “Non-basic”. Finally, “Polarity” classifies a relation as either “Positive” or “Negative”. Combining these four parameters, Sanders et al. suggest a set of 17 relations shown in Table 7. Following the approach in Reference [66], a small scale annotation (DiscAn corpus) of discourse relations is carried out in Reference [65]. As in the previous approaches, the annotation procedure involves pre-segmentation of the text into discourse units and annotation of discourse relations between those units. The taxonomy of relations is extended and modified considerably to accommodate the range of observed examples. For example, the “Basic Operation” primitive is extended to include “Temporal” and “Conditional” relations, although the “Source of Coherence” primitive is not applicable to Temporal relations.

Table 7 — CCR and DiscAn

Relation defined in Reference [66]	Basic Operation	Source of coherence	Order of segments	Polarity
1. Cause-consequence	Causal	Semantic	Basic	Positive
2. Contrastive cause-consequence	Causal	Semantic	Basic	Negative
3. Consequence-cause	Causal	Semantic	NonBasic	Positive
4. Contrastive consequence-cause	Causal	Semantic	NonBasic	Negative
5. Argument-claim	Causal	Pragmatic	Basic	Positive
6. Instrument-goal	Causal	Pragmatic	Basic	Positive
7. Condition-consequence	Causal	Pragmatic	Basic	Positive
8. Contrastive argument-claim	Causal	Pragmatic	Basic	Negative
9. Claim-argument	Causal	Pragmatic	NonBasic	Positive
10. Goal-instrument	Causal	Pragmatic	NonBasic	Positive
11. Consequence-condition	Causal	Pragmatic	NonBasic	Positive
12. Contrastive claim-argument	Causal	Pragmatic	NonBasic	Negative
13. List	Additive	Semantic	NA	Positive
14. Exception	Additive	Semantic	NA	Negative
15. Opposition	Additive	Semantic	NA	Negative
16. Enumeration	Additive	Pragmatic	NA	Positive
17. Concession	Additive	Pragmatic	NA	Negative

6.8 Penn Discourse Treebank (PDTB)

The annotation framework of the PDTB[59][61] is driven by the DLTAG[16][75] conception of discourse relations as predicates in the discourse representation. However, further claims in DLTAG about DAGs as data structures for discourse (i.e. permitting multiple parents but not crossed dependencies) or of the separate treatment of structural and anaphoric discourse dependencies are not carried over to the PDTB. The PDTB annotates discourse relations and their arguments while remaining agnostic about discourse structure representation and non-committal about the distinction between structural and anaphoric dependencies.

In contrast to all other approaches to discourse relation annotation, the texts in the PDTB are not pre-segmented before the annotation of the relations. Rather, the annotation involves identification of the explicit and implicit triggers of relations first, and is then followed by identification of the arguments and the relation semantics. Two kinds of triggers for relations are considered: (a) explicit connectives from well-defined classes whose arguments can be non-adjacent, and (b) adjacency between sentences (including certain types of intra-sentential clauses), which can lead to the recognition of either an

implicit relation (for which a connective is inserted as part of the annotation), an explicit relation realized with some non-connective expression (AltLex, short for “Alternative Lexicalizations”), an entity-based coherence relation (EntRel), or non-existence of any kind of relation (NoRel). With the second kind of trigger, arguments of the relation are constrained to include the two adjacent sentences, although the extent of the argument can be reduced to less than a sentence or extended to include additional sentences to yield complex discourse units as arguments. In general, there are no syntactic constraints to how far an argument can extend (although some constraints are defined for EntRel and NoRel). Thus, arguments can be single clauses, sentences, or multiple clauses or sentences. From a semantic point of view, however, an argument must contain the *minimal* amount of text that is required for interpreting the relation. Arguments are simply labelled Arg1 and Arg2 based on syntactic criteria: for explicit connectives, Arg2 is the argument to which the connective is syntactically bound, while Arg1 is the other argument; and for relations marked between adjacent sentences, Arg1 and Arg2 naming follows their textual order. The relation classification in the PDTB is a hierarchical classification, shown in Figure 2. Semantic annotation of relations allows for (a) back-off to a higher level in the hierarchy when more fine-grained distinctions are difficult to make, and (b) annotation of multiple relations. EntRel, while being a semantic relation, is not considered a discourse relation with situations as arguments, and is therefore not included in the classification.




Figure 2 — PDTB classification of discourse relations

6.9 Mapping of DR-core discourse relations to existing classifications

Tables 8 to 11 show the correspondence between the DR-core relations and the relations seen in the various classifications discussed in this document. The mapping shown is based on a comparison of the relation definitions provided in the various frameworks. From RST, a few of the presentational relations have also been included since they cover the same kinds of examples, although it may be noted that the presentational type of meaning is described in RST to capture speaker intentions (i.e. speaker's belief of the intended effect on the hearer), so the correspondence with these presentational relations in RST is not strict. On the other hand, it may be possible to view this subset of the presentational relations as also subject-matter relations.

Table 8 — Mapping between DR-core relations and RST/RST Treebank

DR-core	RST	RST Treebank
Cause	<i>Vol. cause, Non-vol. cause, Vol. result, Non-vol. result, Evidence, Justify</i>	<i>Cause, Consequence, Result, Evidence, Explanation-argumentative, Reason</i>
Condition	<i>Condition</i>	<i>Condition, Contingency, Hypothetical</i>
Negative condition	<i>Otherwise</i>	<i>Otherwise</i>
Purpose	<i>Purpose</i>	<i>Purpose</i>
Manner	—	<i>Manner, Means</i>
Concession	<i>Concession</i>	<i>Antithesis, Concession, Preference</i>
Contrast	<i>Contrast</i>	<i>Comparison</i>
Exception	—	—
Similarity	—	<i>Analogy, Proportion</i>
Substitution	<i>Antithesis</i>	—
Conjunction	<i>Joint</i>	<i>List</i>
Disjunction	<i>Joint</i>	<i>Disjunction</i>
Exemplification	<i>Elaboration (set-member)</i>	<i>Elaboration-set-member, Example</i>
Elaboration	<i>Elaboration (all others), Summary</i>	<i>Conclusion, Elaboration-general-specific, Elaboration-part-whole, Elaboration-process-step, Summary</i>
Restatement	<i>Restatement</i>	—

Table 9 — Mapping between DR-core relations and HTDC/GraphBank

DR-core	HTDC	GraphBank
Cause	<i>Explanation, Cause, Evaluation</i>	<i>Cause-effect, Elaboration</i>
Condition	—	<i>Condition</i>
Negative condition	—	—
Purpose	—	—
Manner	—	—
Concession	<i>Violated expectation</i>	<i>Violated expectation</i>
Contrast	<i>Contrast</i>	<i>Contrast</i>
Exception	—	—
Similarity	<i>Parallel</i>	<i>Similarity</i>
Substitution	—	—
Conjunction	<i>Parallel</i>	<i>Similarity</i>
Disjunction	—	—
Exemplification	<i>Exemplification, Generalization</i>	<i>Example, Generalization</i>

Table 9 (continued)

DR-core	HTDC	GraphBank
Elaboration	<i>Elaboration</i>	<i>Elaboration</i>
Restatement	<i>Elaboration</i>	<i>Elaboration</i>
Synchrony	—	
Asynchrony	<i>Occasion</i>	<i>Temporal Sequence</i>
Expansion	<i>Background</i>	<i>Elaboration</i>
Unmapped	—	<i>Attribution, Same-segment</i>

Table 10 — Mapping between DR-core relations, PDTB and CCR

DR-core	PDTB	CCR/DiscAn
Cause	Reason, Result, Justification	Causal-Semantic-Basic-Positive, Causal-Semantic-NonBasic-Positive, Causal-Pragmatic-Basic-Positive, Causal-Pragmatic-NonBasic-Positive
Condition	Hypothetical, General, Unreal Past, Unreal Present, Factual Past, Factual Present	Causal-Semantic-Basic-Positive, Causal-Semantic-NonBasic-Positive, Causal-Pragmatic-Basic-Positive, Causal-Pragmatic-NonBasic-Positive
Negative condition	Condition	—
Purpose	Result	Causal-Pragmatic-Basic-Positive, Causal-Pragmatic-NonBasic-Positive
Manner	—	Additive-Semantic-Basic-Positive, Additive-Semantic-NonBasic-Positive
Concession	Expectation, Contra-expectation	Causal-Semantic-Basic-Negative, Causal-Semantic-NonBasic-Negative
Contrast	Juxtaposition, Opposition	Additive-Semantic-Negative
Exception	Exception	Additive-Semantic-Negative
Similarity	Conjunction	Additive-Semantic-Positive
Substitution	Chosen alternative	Additive-Semantic-Negative
Conjunction	Conjunction, List	Additive-Semantic-Positive
Disjunction	Disjunctive, Conjunctive	Additive-Semantic-Negative
Exemplification	Instantiation	Additive-Semantic-Positive
Elaboration	Generalization, Specification,	Additive-Semantic-Positive
Restatement	Equivalence	—
Synchrony	Synchronous	—

Table 10 (continued)

DR-core	PDTB	CCR/DiscAn
Asynchrony	Precedence, Succession	—
Expansion	EntRel	Additive-Semantic-Positive
Unmapped	—	—

Table 11 — Mapping between DR-core Relations and SDRT/DISCOR/ANNODIS

DR-core	SDRT	DISCOR	ANNODIS
Cause	<i>Explanation, Result</i>	<i>Explanation, Result</i>	<i>Explanation, Result</i>
Condition	<i>Consequence</i>	<i>Consequence</i>	<i>Conditional</i>
Negative condition	<i>Consequence</i>	<i>Consequence</i>	<i>Conditional</i>
Purpose	<i>Explanation</i>	<i>Explanation</i>	<i>Goal</i>
Manner	<i>Elaboration</i>	<i>Elaboration</i>	<i>Elaboration</i>
Concession	<i>Contrast</i>	<i>Contrast</i>	<i>Contrast</i>
Contrast	<i>Contrast</i>	<i>Contrast</i>	<i>Contrast</i>
Exception	—	—	—
Similarity	<i>Parallel</i>	<i>Parallel</i>	<i>Parallel</i>
Substitution	—	—	—
Conjunction	<i>Continuation</i>	<i>Continuation</i>	<i>Continuation</i>
Disjunction	<i>Alternation</i>	<i>Alternation</i>	<i>Alternation</i>
Exemplification	<i>Elaboration</i>	<i>Elaboration</i>	<i>Elaboration</i>
Elaboration	<i>Elaboration</i>	<i>Elaboration</i>	<i>Elaboration</i>
Restatement	<i>Elaboration</i>	<i>Elaboration</i>	<i>Elaboration</i>
Synchrony	—	—	—
Asynchrony	<i>Narration</i>	<i>Narration, Precondition</i>	<i>Narration, Flashback</i>
Expansion	<i>Background, Elaboration</i>	<i>Background, Elaboration, Commentary</i>	<i>Background, Entity-Elaboration, Comment</i>
Unmapped	—	<i>Attribution, Source</i>	<i>Attribution, Frame, Temporal-location</i>

In some more recent works, there have been similar efforts^{[67][81]} to relate and map discourse relations across different frameworks. Reference [81] proposes a new taxonomy of relations organized in four top level-classes and detail a mapping between existing annotations, although this is currently limited to only the RST and SDRT frameworks. The proposal in Reference [67] is also for unifying discourse relations across frameworks in terms of a limited set of dimensions, along the lines of the dimensions in the CCR framework, and using them for a mapping between relations in PDTB, RST and SDRT.

7 Interactions of this document with other annotation schemes

7.1 Overlapping annotation schemes

Some of the relations defined in this document occur also in other ISO standards for semantic annotation; for instance, as a relation between an event and one of its participants (i.e. as a semantic role as defined in ISO 24617-4), or as a relation between an event and its time of occurrence (i.e. as a temporal relation as defined in ISO 24617-1). Moreover, some discourse relations are specific for *interactive* discourse (dialogue), and have been defined in ISO 24617-2. DR-core thus overlaps with the schemes for annotating semantic roles, time and events, and dialogue acts. Since overlaps between

annotation schemes are a potential source of semantic and terminological inconsistencies (see the discussion on overlapping schemes in ISO 24617-6), the following subclauses discuss the overlaps between a DR-core on the one hand, and the ISO standards for semantic role annotation (7.2), for the annotation of time and events (7.3), and for dialogue acts (7.4).

7.2 Discourse relations and semantic roles

The following discourse relations, defined in [Clause 5](#), are also defined as semantic roles in ISO 24617-4: Cause, Purpose and Manner. For each of these roles, the question arises as to whether the two annotation schemes define the same concept.

The semantic role Cause is defined in ISO 24617-4 as: “Participant in an event that initiates the event, but does not act with any intentionality or consciousness; the participant exists independently of the event.” In [Clause 5](#), the Cause relation is defined as a relation between two situations where the occurrence of one can explain the occurrence of the other. These definitions do not define exactly the same relation, since in the former (1) of the arguments can be an entity other than a situation, as in “Santa frightened the kids”; and (2) if two eventualities are related, one of them is a participant in the other, which is not the case in the latter. On the other hand, the conditions that define a Cause discourse relation between two events do also hold for the Cause semantic role. Cause as a discourse relation is thus more general than Cause as a semantic role.

The semantic role Purpose is defined in ISO 24617-4 as “Set of facts or circumstances that an agent wishes or intends to accomplish by performing some intentional action.” As a discourse relation, Purpose is defined as a relation between two situations A and B, such that A enables B.” The notion of a “situation” is broad enough to encompass a “set of facts of circumstances”, and “enables” is more general than “wishing or intending to accomplish”, so the discourse relation Purpose is more general than the semantic role of the same name.

The semantic role Manner is defined in ISO 24617-4 as: “The way or style of performing an action or the degree/strength of a cognitive or emotional state.” As a discourse relation, Manner is defined as a relation between two situations A and B where A contains information about how B comes about or occurs. These definitions are identical except that the definition of Manner as a semantic role characterizes the argument B as an “action or a cognitive or emotional state”, whereas in the definition as a discourse relation B is something that “comes about or occurs”, so the difference is between B being possibly a “cognitive or emotional state” or a more general situation. Since it hardly seems possible to describe the manner of occurrence of a situation other than as a degree or strength of a cognitive or emotional state (for example, it is not possible to describe the manner of a negated event), the two definitions can be considered to in fact define the same concept.

It may be concluded that the overlap between the DR-core scheme of this document and ISO 24617-4 is harmless.

7.3 Discourse relations and temporal relations

The following discourse relations, defined in [Clause 5](#), are of a temporal nature and correspond to temporal relations that occur in TimeML.

- a) Synchrony, which corresponds to the TimeML relations Simultaneous and During.
- b) Asynchrony, which corresponds to the TimeML relations After and Before.

The temporal relations in TimeML are based on Allen’s interval calculus, which distinguishes 13 relations between temporal intervals: (1 and 2) before/after; (3 and 4) overlaps/overlapped_by; (5 and 6) start/started_by; (7 and 8) finish/finished_by; (9 and 10) during/contains; (11 and 12) meet/met_by; and (13) equality. Six of these are the inverse of another relation, differing only in the order of the arguments, and are strictly speaking redundant, so there are seven primitive relations.

TimeML defines 12 relations, six of which are either relations between events, or between an event and a temporal object (interval or time point), or between two temporal objects; the other six apply only to two *events*. Five relations are the inverse of another relation, differing only in the order of the

arguments, and are strictly speaking redundant, so there are seven primitive relations. The fact that some of the temporal relations in TimeML can relate two events, an event and a time, or two times, sometimes causes some conceptual unclarity, which the temporal relations defined in DR-core do not have. Consider, for example, the relation *Simultaneous*, defined in TimeML as follows:

Simultaneous: Two events are judged simultaneous if they happen at the same time, or are temporally indistinguishable in context, i.e. they occur close enough so that further distinguishing their times makes no difference to the temporal interpretation of the text. This is also used for expressing the duration of an ongoing event, as in *Mary taught from 2 to 4*.

Conceptually, saying that two events co-occur seems something rather different from saying when an event occurred. From a terminological point of view it is commendable to keep the term *Synchrony*, defined in DR-core as that of simultaneous occurrence of two events, and as such coinciding with the inter-event interpretation of the *Simultaneous* relation of TimeML.

Table 12 shows the correspondences between the temporal discourse relations defined in DR-core, the temporal relations defined in TimeML, and those defined in Allen's calculus. The temporal discourse relations defined in DR-core are clearly less specific than those defined in TimeML.

Table 12 — Temporal relations of DR-core, TimeML and Allen's temporal calculus

DR-core	TimeML	Allen's calculus
Synchrony	Simultaneous (between events)	Equality
Asynchrony	Before, After (between events)	Before, After
Asynchrony	I_Before, I_After (between events)	Meet, Met_By
Synchrony	—	Overlap
Synchrony	Is_Included, Includes (between events)	—
Synchrony	During	During, Contains
—	Begins, Begun_By (between events)	Starts, Started_By
—	Ends, Ended_By (between events)	Finishes, Finished_By

7.4 Discourse relations and semantic relations between dialogue acts

Due to the interactive nature of a dialogue, utterances in a dialogue can be related by other kinds of relations than those that connect clauses in a written text. ISO 24617-2 distinguishes two such relations, *functional dependences* and *feedback dependences*.

Functional dependences occur due to the fact that some types of dialogue act are inherently responsive in nature; they presuppose the occurrence of certain other dialogue acts. For example, an Answer presupposes a Question; a Reject Offer presupposes an Offer; an Accept Apology presupposes an Apology, and so on. The meaning of a responsive dialogue act depends on that of the dialogue act that it responds to. This is obvious for short answers like "Yes" and "No", "at 4 p.m.", and for a short Accept Offer like "Sugar please", but it is also true for seemingly full-fledged answers such as Example 21, whose meaning depends on whether it is an answer to Example 22a or to Example 22b:

Example 21 I'm expecting Jamie, Andrea and Elsa to come tonight.

Example 22 a) Who do you expect to be there tonight?

b) Which of the girls from the office do you expect to be there?

The meaning of an answer depends in general on the meaning of the question that it answers, and similarly for responses to requests, offers, suggestions, and other responsive dialogue acts.

Feedback dependences occur whenever an explicit feedback act is performed, as in Example 23:

Example 23 A: Continue to go south until you see a haystack to your left.

B: Okay.

This is an example of an Instruct act followed by a positive feedback act. Feedback in communication is the phenomenon that a dialogue participant provides or elicits information about the processing of previous utterances. To understand the meaning of a feedback act one must know which previous utterance(s) the feedback is about. In Example 23, this is the immediately preceding utterance, which is often the case, but things may be more complicated, as in the following dialogue fragment (from the HCRC Map Task corpus):

Example 24 1) A: Keep going, slightly east, until you come to the old mill.

2) B: Mmm, I don't see a mill.

3) A: To your right, when you pass the haystack, do you see a barn?

4) B: Yes.

5) A: And a little further down, you see the ruins of a mill?

6) B: Ah okay, those ruins are from a mill.

With utterance 6), B provides positive feedback about both A's utterance 5) and A's first utterance. Positive feedback (reporting successful processing) is mostly related either to the immediately preceding utterance of the addressee or to a sequence of the addressee's utterances; negative feedback is nearly always related to the immediately preceding utterance, as in the case of utterance 2). All in all, the meaning of a feedback act evidently depends on the utterances about which it provides or elicits feedback; this dependence is called a "feedback dependence".

To make the present core annotation scheme applicable to interactive discourse and consistent with ISO 24617-2, the set of discourse relations include functional dependence and feedback dependence. It may be noted that both dependence relations, while specific for interactive discourse, do appear in written fictional texts, such as novels, when the personages involved are engaged in dialogue.

8 DRelML: Discourse Relations Markup Language

8.1 Overview

The Discourse Relations Markup Language (DRelML) has been designed in accordance with the SemAF Principles of semantic annotation (ISO 24617-6), which implements the distinction between *annotations* and *representations* made in the Linguistic Annotation Framework³⁾ (ISO 24612). This means that the definition of an annotation language consists of three parts:

- a) an abstract syntax, which specifies the class of annotation structures;
- b) a formal semantics, describing the meaning of the annotation structures defined by the abstract syntax;
- c) a concrete syntax, specifying a reference format for representing the annotation structures defined by the abstract syntax.

The concrete syntax is required to be *complete* and *unambiguous* relative to the abstract syntax, i.e. (a) the concrete syntax defines a representation for every structure defined by the abstract syntax; and (b) every expression defined by the concrete syntax represents one and only one structure defined by the abstract syntax. A concrete syntax with these two properties defines an *ideal representation format*. Any ideal representation format can be converted through a meaning-preserving mapping to

3) The term "annotation" refers to the linguistic information that is added to regions of primary data, independent of the format in which the information is represented; "representation" refers to the format in which an annotation is rendered.

any other ideal representation format.⁴⁾ Worth noting is also that a mapping strategy can be defined to convert from an abstract syntax to a representation in GrAF⁵⁾ format,^[24] as shown in Reference [22]. In addition to allowing for annotation schemes to be represented uniformly across languages, domains, and genres, this may be useful for combining PDTB annotations with GrAF renderings of PropBank and other annotations that have been done on the WSJ, including syntactic (PTB) annotations.

It may be worth pointing out that annotators who use DRelML only deal with concrete XML representations; they do not have to worry about the underlying abstract syntax. For the sake of interoperability with alternative representation formats, the semantics of concrete DRelML representations is defined by the semantics of the annotation structures that they encode. The extent to which alternative representations encode the same underlying abstract annotation structures determines the extent of their interoperability. The DRelML semantics is, in practice, mainly of interest for computational applications that exploit the semantic information in annotation representations.

8.2 DRelML abstract syntax and semantics

The abstract syntax of DRelML consists of a “conceptual inventory”, specifying the concepts from which annotations are built up, and a specification of the possible ways of combining these elements to form annotation structures. The conceptual inventory specifies the ingredients of discourse relation annotations according to the metamodel shown in Figure 1. This includes the set of discourse relations defined in Clause 5; the argument roles defined for each of these relations; and the set of types that the arguments may have.

The annotation of a discourse relation includes the identification of the regions of primary data corresponding to the arguments of the relation, and in the case of an explicit discourse relation also the region where the relation is expressed. In stand-off annotation, this is done by means of “markables”, elements such as word tokens or morphosyntactic units (typically as the output of a tokenizer or a morphosyntactic analyser) which point to regions of primary data. The conceptual inventory therefore includes also a set of markables, which is specific for a particular annotation task.⁶⁾ Altogether, the DRelML conceptual inventory consists of the following specifications:

- a) D , a set of discourse relations;
- b) R , a set of argument roles for discourse relations;
- c) A function α from D to $R \times R$, which assigns a pair of argument roles to each discourse relation;
- d) M , a set of markables that identify the segments of primary data to be marked up in a given annotation task;
- e) T , a set of argument types (including eventualities, as defined in ISO 24617-1, and dialogue acts, as defined in ISO 24617-2).

An annotation structure is a set of *entity structures*, which contain semantic information about a region of primary data, and *link structures*, which specify a semantic relation between two such regions.

Formally, an entity structure is one of the following:

- a relation entity structure, which is a pair $\langle m, r_j \rangle$ consisting of a markable m (a member of M) and a discourse relation r_j (a member of the set D);

4) See References [6] and [8] for formal definitions and proofs.

5) GrAF may be considered as a pivot format into which well-formed annotation schemes may be mapped, thus guaranteeing syntactic consistency and completeness for the purposes of comparison, merging, and transduction to other formats.

6) If the annotation task consists of marking up pre-segmented data, then the set of markables is given at the start of the task. If, by contrast, unsegmented data are to be annotated, then the set of markables is constructed as part of the task.