

2 Theoretical Models of Discourse Relations

2.1 INTRODUCTION

During the last decades, having an explicit representation of discourse structure has become a pressing need for many applications in computational linguistics, such as automatic summarization and human–computer interactions. During the same period, the development of new theories and methodologies in psycholinguistics has also meant that the study of language processing could go beyond the level of isolated sentences. This evolution also implied the need for a cognitively motivated representation of discourse structure, accounting for discourse coherence. In this context, theoretical models of discourse structure have started to emerge since the 1980s in order to provide such explicit representations. We will present a selection of the most prominent ones in this chapter, starting with Rhetorical Structure Theory (2.2) and Segmented Discourse Representation Theory (2.3), two models that share the goal of providing a global representation of discourse structure encompassing every segment of a text, thus going beyond the level of local discourse relations. We will then present lexically grounded approaches to discourse structure that anchor the study of discourse relations in the use of connectives, in particular the Penn Discourse Treebank project (2.4). We will finally present the Cognitive Approach to Coherence Relations, a model emphasizing the need to provide a cognitively plausible account of discourse relations, in the form of a set of cognitively motivated primitives into which all of them can be decomposed (2.5).

In this chapter, our main objective is to provide a succinct description of each model, emphasizing their main goals, and discussing their advantages and limitations. We will also list their specificities compared to other models, and analyze the main differences between them. We will focus more specifically on the aspects of these models that have to do with the description of discourse

relations, and leave aside other components linked to global discourse structure such as schemas, as well as the question of discourse segmentation. We refer interested readers to relevant publications about these aspects of discourse structure at the end of the chapter. For each model, we will present the type of research to which it has been applied, and the data that have been produced in the form of annotated corpora. As we will see, all these models have been used to annotate large corpora with discourse relations. An important issue is therefore to establish mappings between the relations annotated in each of them, in order to compare data from one corpus to the others. We discuss various options for comparing annotations across models in the last section (2.6). A specific model developed for the annotation of discourse markers (and connectives) in spoken discourse focusing on their polyfunctionality and polysemy will be presented in Chapter 3.

2.2 RHETORICAL STRUCTURE THEORY

Rhetorical Structure Theory (RST) was one of the first models developed in the 1980s (Mann & Thompson, 1988) as an attempt to provide a global theory of discourse structure. The initial goal was to provide a tool that could be used in computer-based text generation. Since then, the theory has also become a valuable descriptive tool in itself, with many different applications.

The starting point for RST comes from the intuitive observation that texts are not made of arbitrary collections of sentences, but rather exhibit internal structure that make them appear coherent to a reader. Thus, in this model, coherence can be defined as an absence of non-sequiturs, in other words clauses following each other without obvious logical links between them. Even though there is no formal obligation that every part of the text is included in an RST analysis, well-formed texts do not usually require not having any elements left out. When performing an RST analysis of a text, the analyst starts by segmenting the text into spans, and then determining the relations between them, called rhetorical relations, a notion similar to the term of discourse relation that we use in this book. Additionally, relations are hierarchical, depending on the lengths of the text spans that they unite. In other words, local relations can be embedded into more global ones within a text.

The list of relations included in RST varies somewhat from author to author, a fact that Mann and Thompson (1988) had already anticipated,

Table 2.1 *List of relations in Mann and Thompson (1988)*

Circumstance	Antithesis and concession
Solutionhood	Antithesis
Elaboration	Concession
Background	Condition and otherwise
Enablement and motivation	Condition
Enablement	Otherwise
Motivation	Interpretation and evaluation
Evidence and justify	Interpretation
Evidence	Evaluation
Justify	Restatement and summary
Relations of cause	Restatement
Volitional cause	Summary
Non-volitional cause	Other relations
Volitional result	Sequence
Non-volitional result	Contrast
Purpose	

as they foresaw that different relations might be needed for different languages or text types (see also Taboada & Mann, 2006a). One of the most widely accepted versions (sometimes called “classical RST”) of the list comes from their 1988 paper, and includes 23 relations, summarized in Table 2.1.

Later on, other relations such as *list*, *means*, *preparation*, *unconditional* and *unless* were added to the list (Mann, 2005). Even though there is no upper limit to the number of relations that can be included in RST, many authors warn against adding a great variety of relations that could not be identified reliably by analysts.

The list of relations proposed by Mann and Thompson (1988) is not really organized as a taxonomy with different families of relations. The authors explain that in their view, there isn’t one single taxonomy that would be entirely appropriate, as different groupings could be made depending on the research question. Still, a division that has often been suggested within RST concerns the opposition between relations that deal with the subject matter of the text, such as elaboration, solutionhood, all types of causal relations, etc. Their function is to be recognized and understood by the reader. The other category includes relations that deal with presentational aspects of the text. In other words, their role is to produce an effect on the reader. For instance, relations of justification are inferred when one segment increases the likelihood that the reader will accept the claim presented in the other segment.

A specificity of RST compared to other frameworks comes from the identification of two different parts for most relations, called the nucleus and the satellite. Nuclei represent the most important part of the relation, whereas satellites are more secondary. If all nuclei are removed from a text, its content is not interpretable anymore. But if satellites are removed, the text, even though incomplete and agrammatical, can still be understood. For example, in a relation of evidence as in (1), the two related text spans include a claim ('nobody is at home') and the evidence backing it up ('the lights are out'). In this case, the claim is the most important span, hence the nucleus, and the evidence is the satellite. To take another example, in a relation of elaboration as in (2), the nucleus contains the basic or main information ('Paul had a great holiday') and the satellite the additional information ('He went swimming...').

- (1) The lights are out, so nobody is at home.
- (2) Paul had a great holiday. He went swimming, ate good food and partied every night.

[constructed examples]

While there is no fixed order for nuclei and satellites within a text, preferential patterns have been observed for some relations. For example, in relations of elaboration, restatement or enablement, the nucleus usually comes first in the text, followed by a satellite. In contrast, for relations of concession, condition or background, the satellite typically comes first and the nucleus second. Other relations do not display a preferential order. Other relations still don't have a segment that is more important than the other, as for example the relation of contrast in (3). The two parts are simply the two sides of the contrast. These relations are called multinuclear. Other examples of multinuclear relations include lists and sequences.

- (3) Helen is blond but Sandra is a brunette.

[constructed example]

Both types of relations can apply either at the sub-sentential level, as in (4) for multinuclear relations and (5) for nuclear relations, or between sentences, as in (6) for multinuclear relations or (7) for nuclear ones.

- (4) Peel the carrots, and slice them into thin slices.
- (5) Good as it may look, I won't eat dessert.
- (6) Peel the carrots, slice them into thin slices. Cook them briefly in the pan, and serve hot.

- (7) I could say a lot more about this topic. But time is up, and
I will stop there.

[constructed examples]

In RST, each relation comes with four types of constraints: a constraint on the nucleus; a constraint on the satellite; a constraint on the combination of the nucleus and the satellite; a constraint on the effect produced. For instance, the relation of evidence poses a constraint on the nucleus that the reader may not believe the content of the nucleus to a degree that is satisfactory to the writer. The constraint on the satellite is that the reader either believes it or will find it credible. The constraint on the relation between the nucleus and satellite states that the reader's comprehension of the satellite will increase their belief in the nucleus. The intended effect is that the reader's belief in the nucleus is increased. A label of evidence can be used only if the analyst is convinced that the writer wanted the effect to be inferred.

The constraints vary greatly from relation to relation. To take another example, the relation of justification does not pose any constraint on the nucleus or the satellite, only on their combination: the fact that the reader comprehends the satellite will increase their willingness to accept the writer's right to present the nucleus. The effect is therefore to increase the reader's willingness to accept the writer's right to present the nucleus.

A big corpus of texts in English, comprising 385 documents from the Penn Treebank made of articles from the Wall Street Journal, has been annotated with RST relations (Carlson & Marcu, 2001). This corpus was built in order to help the development of computer-based applications such as text summarization, machine translation and document retrieval (Taboada & Mann, 2006b). It has also been very useful for answering important questions about the description of discourse relations, as well as their signaling. As the examples above illustrate, RST is focused on defining the way text spans are connected by different types of relations. Since all text spans are included in the annotation, it becomes possible to compare the ways in which various relations are signaled.

Das and Taboada (2018) annotated all the signaling devices used to convey relations in the RST corpus, including connectives but also other lexical, syntactic, semantic, graphical and genre features, and found that 90 percent of relations in their data were signaled, sometimes even with multiple signals. However, connectives represented a small portion of the signaling devices, as only 11 percent of the relations were signaled exclusively by connectives, against 75 percent of

the relations that were signaled by other means. This explains why connectives are somewhat marginal for RST, a theory mostly interested in the segmentation of texts into discourse relations. On the one side, connectives are not needed to convey a relation, and when they are used, they do not unambiguously convey a relation because many of them are polyfunctional (see Chapter 3). Yet, they also found important differences between relations that were very rarely marked, such as background and restatement, and relations that were very often explicitly marked like concession and condition. These observations have in turn been quite useful to study the role of connectives and other signaling devices for the processing and acquisition of discourse relations (see Chapters 6 and 8).

Since its conception, RST has been applied for analyses in many different domains (Taboada & Mann, 2006b), in addition to the computer applications mentioned above. Even though it was originally designed with English in mind, RST has been used to compare the communication of discourse relations between different languages (see Chapter 7). It has also been used beyond the analysis of monologic texts, and applied to dialogues (e.g., Daradoumis, 1996), and even to studying the links between speech and gestures in communication (de Carolis et al., 2000). RST has also been used to compare the communication of relations across various genres such as academic (Benwell, 1999) and argumentative (Azar, 1999) discourse. Finally, RST has been used to evaluate text writing in L1 (Bouwer, 1998) and L2 (Kong, 1998).

To summarize, RST provides a way to analyze the structure of texts by decomposing them into discourse relations. For this reason, this theory is centered on the notion of discourse coherence and offers a limited place for the study of connectives.

2.3 SEGMENTED DISCOURSE REPRESENTATION THEORY

Segmented Discourse Representation Theory (SDRT) was developed in the 1990s (Asher, 1993; Lascarides & Asher, 1993) based on two different trends of research from the 1980s, encompassing both formal semantics and theories of discourse. It is first based on Discourse Representation Theory (Kamp & Reyle, 1993), a formal semantics model developed in order to account for discourse phenomena going beyond the level of individual sentences such as anaphoric relations, and second on theories of discourse structure with applications to computational linguistics, such as Rhetorical Structure Theory (see Section 2.1) and Centering Theory (Grosz & Sidner, 1986). SDRT aims at keeping the

formal rigor of DRT while using the notion of discourse relation in order to resolve problems arising from this framework.

One such problem is related to the temporal interpretation of discourse, as illustrated with the following pair of examples:

(8) Max opened the door. The room was pitch dark.

(9) Max switched off the light. The room was pitch dark.

[from Lascarides & Asher, 1993: 437]

In both examples, sentences contain verbs in the past tense, and have a similar grammatical structure. In addition, the first sentence describes a punctual event whereas the second one describes a stable state of affairs. Yet, there is a major difference between them in terms of process. While in (8) the state of darkness covers the whole event of opening the door, in (9) this same state of darkness happens only after Max has switched off the light. Yet, DRT does not provide any means to differentiate between these two situations, because it does not take into account the type of discourse relation holding the two sentences together. In (8), the relation is one of background, but in (9) it is a relation of result. With the information from the discourse relation, the distinction between the two examples becomes clear. SDRT was precisely developed to provide means to incorporate discourse relations into a logical representation of discourse, thus accounting for these differences.

A specificity of SDRT compared to other discourse models is its integration of two types of analyses of discourse structures. First, a bottom-up approach starting from minimal discourse units and linking them with discourse relations in a recursive fashion (relations can be embedded in one another). Second, a top-down construction starting from a full or partial discourse structure and identifying signals of global text organization. While Asher et al. (2017) emphasize that both types of structures can lead to similar results, they also note that analyses taking one or the other approach typically focus on different aspects of discourse structure: local relations on the one side, and more global structures on the other. These two analyses also involve a different focus from a cognitive perspective. Conducting a top-down analysis means that readers are believed to look for global textual coherence before assigning local links between sentences, and to focus first on global structures such as thematic continuity or discontinuity rather than more local discourse relations. Within the SDRT framework, both types of structures are deemed important and complementary, and both have been annotated in corpus data (see below).

Table 2.2 *Discourse relations from SDRT (Reese et al., 2007: 8)*

Coordinating relations		Subordinating relations	
Veridical	Nonveridical	Veridical	Nonveridical
Continuation	Consequence	Background	Attribution
Narration	Alternation	Elaboration	
Result		Explanation	
Contrast		Commentary	
Parallel		Source	
Precondition			

At the level of discourse relations, SDRT takes a medial position between theories that make use of a high number of relations (such as RST and PDTB), and more minimalist models like the two relations used by Grosz and Sidner (1986). In total, 14 relations have been selected to account for written texts, which was the original objective. These relations are listed in Table 2.2. They are classified first based on a grammatical criterion: whether they introduce horizontal relations between coordinated segments, or whether they introduce a hierarchic relation with a subordinate clause. It also distinguishes between veridical relations that entail the content of their arguments, and nonveridical relations that do not entail the content of at least one of the arguments.

In SDRT, discourse relations are characterized semantically. With this precise semantic description, it can be verified whether two relations are the same, if one of them entails the other or if they are incompatible. An example of such a definition is given below for the relation of explanation, taken from Reese et al. (2007: 12):

When α and β introduce eventualities in the dynamic sense (i.e. existential quantification over eventualities occurs with wide scope over modal operators, negation or non-existential quantifiers), $\text{Explanation}(\alpha, \beta)$ holds when the main eventuality of β is understood as the cause of the eventuality in α . *Explanation* has temporal consequences, viz. that the eventuality described in β precedes (or overlaps) the eventuality described by α . *Because* is a monotonic cue for *Explanation* [...] ‘After’ and ‘when’ sometimes signal *Explanation*.

This example illustrates several important points. First, relations are described in relation to one another, as for example the link between temporality and explanation. Second, relations are defined independently of discourse connectives and other markers that are deemed to be

too ambiguous, but they are still listed as potential indicators of a relation, such as *because*, *after* and *when* in the case of explanation relations. The strength of each signal is also indicated. In that sense, SDRT is similar to RST, a theory also focusing on relations rather than the signals conveying them. Another similarity between these two models is that all textual segments are included in the analysis, as all of them except for the first segment in the text are hypothesized to be linked to at least another with a discourse relation in coherent texts.

There are also a couple of important differences between RST and SDRT, in addition to the different number of relations included in each model, and the way they are labelled. Contrary to RST, in SDRT several relations can be presented simultaneously between two discourse segments, as for example in (10). In this example, the two segments are linked both by a relation of contrast and of narration, as the presence of the two connectives *but* and *then* indicates. If only one relation were allowed between these segments as in RST, analysts would have to choose between them, and thus lose part of the information conveyed.

(10) John gave Mary a book, but then he took it back.

[translated from Busquets, Vieu & Asher, 2001: 82]

There is also another difference in the representation of discourse relations between the two models. In SDRT, the structure of discourse takes the form of graphs rather than trees. This implies the possibility of having attachments between parts of texts that are not contiguous, as in (11) from Asher et al. (2017: 1245) where the two discourse segments 31 and 33 are linked by a relation of contrast, even though they are not contiguous.

(11) [In 1988, Kidder eked out a \$ 46 million profit,]₃₁ [mainly because of severe cost cutting.]₃₂ [Its 1,400-member brokerage operation reported an estimated \$ 5 million loss last year,]₃₃ [although Kidder expects to turn a profit this year]₃₄.

[RST Treebank, wsj_0604]

Similarly to RST, in the PDTB framework, implicit relations are annotated only between adjacent segments (see Section 2.4). As a result, these long-distance attachments cannot be represented in these models. This means that part of the existing relations within a text are missed.

Ten years after the initial version of SDRT, Asher and Lascarides (2003) proposed an extension of the model to clarify or simplify some issues, and to make it more suitable for a broader range of linguistic phenomena. One of the extensions was to include relations for the

annotation of dialogues, as the original SDRT version was conceived for written genres. In order to account for the specificities of dialogues, relations linked to questions and requests were added to the general relation of elaboration. Another addition was the relation of adjacency pairs, in order to account for cases when the first segment is a question and the next one is an answer. Another typical aspect of dialogues is that speakers often correct each other. A relation of correction was thus added to account for this phenomenon.

Another novelty was to make the theory modular, in order to separate various aspects of discourse interpretation. In this view, the inference of discourse relations is a specific module that takes as input underspecified semantic representations, world knowledge and lexical information. This module is deemed to be the one gluing all the other aspects of discourse interpretation together, hence the importance of discourse relations for coherence.

SDRT analyses have been implemented in a large corpus of French written texts: the ANNODIS corpus¹ (Reese et al., 2007; Afantenos et al., 2012). This corpus is multi-genre, as it includes texts from news and encyclopedia articles, linguistics research papers, and international relations reports. It therefore encompasses narrative, expository and argumentative genres (see Chapter 7). The ANNODIS corpus includes two types of annotations. A bottom-up annotation of elementary and complex discourse units linked by discourse relations, and a top-down annotation of high-level structures such as enumerative structures and topical chains, in line with the two types of discourse structures identified in SDRT. As both types of annotations required texts of different length (short for bottom-up annotations and longer for top-down ones), they were performed on distinct subparts of the corpus. The annotation of discourse relations was performed in three phases. In a first phase, two naïve annotators annotated 50 documents, and their input was used to create an annotation manual (Reese et al., 2007) describing the relations and giving information about discourse segmentation. In a second phase, three students double-annotated 86 documents after receiving training. In a last phase, an expert annotator adjudicated and corrected the naïve annotations in order to reach a final version of the corpus.

Thanks to these annotations, ANNODIS is a useful resource to compare the frequency of different relations across genres, and to link the occurrence of discourse relations with other discourse phenomena such as the use of pronouns. Since it uses an onomasiological approach

¹ <http://redac.univ-tlse2.fr/corpus/annodis/>.

to discourse structure, focusing on discourse relations and linking all textual segments with at least one other, this corpus has been useful in showing how relations are realized in discourse and in identifying all the different linguistic forms that can be used to signal each relation, contrary to models that start from markers as a way to identify relations. From its inception, SDRT has been of interest for natural language processing applications, and the ANNODIS corpus has been used as training for systems dealing with discourse structure prediction, discourse parsing, relation labeling and sentiment analysis. The inclusion of dialogic data into SDRT has been tested in another project (STAC) involving the annotation of on-line chat dialogues (Asher & Paul, 2018).

In sum, SDRT is a formal model of discourse structure that includes, but is not limited to, the annotation of discourse relations. Like RST, the focus is placed on discourse relations whereas connectives play a rather marginal role due to their ambiguities (see Chapter 3).

2.4 THE PENN DISCOURSE TREEBANK FRAMEWORK

The Penn Discourse Treebank (PDTB) is a project that started in the first decade of the twenty-first century to annotate the million words corpus from the Wall Street Journal with discourse information² (Webber et al., 2006). More specifically, the idea behind this project is to annotate explicit and implicit discourse connectives throughout the corpus. Implicit connectives correspond to cases where no connective was used in the text, but the annotator judged that there was still a discourse relation that could be inferred between two adjoining text segments, and that this link could be adequately expressed by a connective. In such cases, they inserted this connective, and it counted as an implicit connective. A major difference between the PDTB framework and the other models we discussed so far is that it is not tied to any theory, and does not aim at building global discourse structures beyond the linking of arguments by connectives.

The PDTB can be considered as a lexically grounded approach to discourse (Prasad, Webber & Joshi, 2014), as relations are intrinsically linked to connectives and other signals that can be used to express them. Indeed, annotators are always guided by connectives first (explicit or implicit) before attempting to label the discourse relation that it communicates. One of the main advantages of the PDTB is that it

² <https://catalog.ldc.upenn.edu/LDC2019T05>.

provides the most extensive resource of annotated connectives available. In the 2.0 version of the corpus, released in 2008 (PDTB Research Group, 2008; Prasad et al., 2008), there were 18 thousand explicitly and 16 thousand implicitly signaled relations. In the PDTB-3 version released in 2019, 17 thousand new relations were added. These correspond mostly to intra-sentential relations, for example, between conjoined verb phrases with segments containing free adjuncts or to-infinitives that had not been annotated in the previous version. In addition to explicit and implicit relations, the PDTB also contains an annotation of alternative lexicalizations or AltLex. These annotations were used when a relation was not conveyed by an explicit connective but the addition of a connective between the two segments would still be inappropriate because it would lead to redundancy, as other information in the sentence also signaled the relation, as in (12):

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

[from Prasad, Webber & Joshi, 2017: 1201]

In this example, the relation of causality cannot be made explicit by adding the connective *because*, as it would be redundant with the information conveyed by the expression “a major reason”. This expression corresponds to a case of AltLex in the PDTB. Given that the identification of alternative lexicalizations was limited to the annotation of implicit relations in which a connective could not be inserted, it is clear that this annotation does not cover all cases in which relations are conveyed by other means than connectives, contrary to corpora annotated within the RST and SDRT frameworks, which include all relations of a given type, independently of its marking.

Two other types of annotations for discourse relations were included. First, the tag EntRel was used when the coherence was entity based, in other words, when the second segment was an extension giving more information about an entity described in the first segment, as in (13) where the second segment provides further information about Hale Milgrim:

- (13) Hale Milgrim, 41 years old, senior vice president, marketing at Elecktra Entertainment Inc., was named president of Capitol Records Inc., a unit of this entertainment concern. Mr. Milgrim succeeds David Berman, who resigned last month.

[from Prasad, Webber & Joshi, 2017: 1201]

Finally, when no relation could be perceived between the segments, the tag NoRel was used. When all annotations are put together, the PDTB-3 now contains over 53 thousand tokens of annotated discourse relations (Prasad, Webber & Lee, 2018).

Given the importance of discourse connectives in this framework, a crucial aspect of the project was to define them in such a way as to label as many different relations as possible, while separating them clearly from neighboring classes. In the PDTB-2, connectives were restricted to four well-defined syntactic classes: subordinating conjunctions (*because, when, if, etc.*); coordinating conjunctions (*and, but, or, etc.*); prepositional phrases (*as a result, in comparison, etc.*); and adverbs (*then, instead, yet, etc.*). This list excluded two neighboring classes of lexical markers: cue phrases like *well* and *so* in the sentence initial position that are used for functions like topic shifts rather than for the communication of discourse relations, and discourse markers like *actually* that do not take scope over two arguments. In the PDTB-3 version, the list of connectives has been enlarged to include prepositional subordinators like *for, with, instead of, etc.* that can also complementize for clauses.

The list of discourse relations included in the PDTB takes the form of a hierarchy encompassing three different levels. This list has evolved between the PDTB-2 and PDTB-3 releases. We focus here on the more recent list of relations included in the PDTB-3 in Table 2.3. The list of senses from the PDTB-2 version is described in the annotation manual (PDTB Research Group, 2008) and early experiments with sense annotation are reported in Miltsakaki et al. (2008).

In this list, Level-3 is used only for relations that can have a different directionality. In other words, one or the other argument linked by the connective can take on a specific role, for example, conveying the goal or the cause segment. This is new compared to Version 2 that used this level to make more fine-grained distinctions between sub-types of relations. In the current version, all relation types can be found at Level-2, whereas Level-1 merely categorizes these relations into four main families. A few fine-grained distinctions that could not be annotated reliably or that occurred very rarely were removed from the list (for example, the various subtypes of conditional relations), while a few relations missing from the previous version were added (for example, a relation of *similarity* in the comparison group).

The annotations of connectives were performed using pre-established lists of connectives from the literature, with which annotators proceeded one connective at a time throughout the entire corpus, so that they could benefit from their increasing expertise. A specific tool was developed to perform the annotation in the corpus (Lee et al.,

Table 2.3 *List of relations from the PDTB-3 (Webber et al., 2019)*

Level-1	Level-2	Level-3
Temporal	Synchronous	–
	Asynchronous	Precedence Succession
Contingency	Cause	Reason Result Neg-Result
		Reason+Belief Result+Belief
		Reason+SpeechAct Result+SpeechAct
	Cause+Belief	Arg-1-as-Cond Arg-2-as-Cond
		–
	Cause+SpeechAct	Arg-1-as-NegCond Arg-2-as-NegCond
		–
	Condition	–
	Condition+SpeechAct	–
	Negative Condition	–
	Negative Condition+SpeechAct	Arg-1-as-Goal Arg-2-as-Goal
		Arg-1-as-Denier Arg-2-as-Denier
	Purpose	Arg-2-as-Denier +SpeechAct
Comparison	Concession	–
	Concession+SpeechAct	–
		–
Expansion	Contrast	–
	Similarity	–
	Conjunction	–
	Disjunction	–
	Equivalence	–
	Exception	Arg-1-as-Excpt Arg-2-as-Excpt
		Arg-1-as-Instance Arg-2-as-Instance
	Instantiation	Arg-1-as-Detail Arg-2-as-Detail
		Arg-1-as-Manner Arg-2-as-Manner
	Level-of-detail	Arg-1-as-Subst Arg-2-as-Subst
	Manner	–
	Substitution	–

2016). This tool was also used to compare two annotations and to adjudicate cases of disagreements. When the disagreement concerned a Level-1 tag, the adjudication was done by a group of expert annotators. When the disagreements concerned Level-2 or Level-3, the

relation was automatically labelled at the higher level, for which there was an agreement between annotators.

In the PDTB, the same two discourse segments could be annotated with more than one discourse relation, for example, when the connective had multiple senses (see Chapter 3), or when the relation was implicit and the annotators inferred more than one relation between them. This enabled the authors to identify the most recurrent cases of double relations. In addition to the annotation of connectives, the PDTB corpus also contains an annotation of the arguments related by connectives, and an indication of attribution of the content of the segments, which can be ascribed to beliefs or assertions performed by the writer or by a person that is being talked about in the text. This led the authors to observe that some connectives can involve a use of different attributions between the two segments. In PDTB-3, another type of annotation has been added, to account for cases when argument 1 involves a question and argument 2 provides an answer to that question. Since questions are treated as dialogue acts in the literature (e.g., Bunt et al., 2020) and these sequences cannot be instantiated by a connective, they are not considered as a new discourse relation, but rather as a complementary phenomenon (Prasad, Webber & Lee, 2018).

A similar approach to the PDTB has been adopted to annotate corpora in other languages such as Arabic (Al-Saif & Markert, 2011), Chinese (Zhou & Xue, 2012), Hindi (Kolachina et al., 2012), Turkish (Zeyrek, Demirahin & Sevdik alı, 2013) and Czech (Zikánová et al., 2010). The relations from PDTB-3 have also been used in an ongoing effort to put together lexicons of connectives from different languages (Stede, Scheffler & Mendes, 2019; see Chapter 4).

Since its release in 2008, the PDTB corpus has been used for various language technology applications, such as the automatic annotation of discourse relations (Pitler et al., 2008), and the prelabeling of connectives to improve the output of machine translation systems (Meyer & Popescu-Belis, 2012). The PDTB has also been used to assess cognitive theories about discourse, such as the continuity hypothesis (see Chapter 6) according to which some discourse relations should be conveyed implicitly more often than others. Thanks to the annotation of explicit and implicit connectives in the PDTB, Asr and Demberg (2012b) were able to compare the ratio of implicitness across relations, and to confirm the existence of a continuity constraint empirically.

To summarize, the PDTB takes a radically different approach from previous models, because it is theory neutral and discourse relations are annotated only in relation to connectives, without searching for a more global discourse structure. This lexical view of discourse relations

has enabled researchers to apply a similar method to typologically diverse languages, and to compare their results (Prasad, Webber & Joshi, 2014).

2.5 A COGNITIVE APPROACH TO COHERENCE RELATIONS

The Cognitive Approach to Coherence Relations (CCR) originated in the 1990s (Sanders, Spooren & Noordman, 1992) as an attempt to provide a cognitively motivated set of primitives to account for the basic features of discourse relations. The idea was to go beyond a simple list of discourse relations, and to characterize each of them in terms of four different primitives. This decomposition was meant to account for all possible cases in terms of basic cognitive principles such as causality. It was also meant to account for the polyfunctionality of some connectives, and showing that their various senses shared some elements in their primitives. For example, the fact that a connective like *and* can be used to convey additive or causal relations but never concessive ones can be accounted for by the basic difference of polarity between positive and negative relations (see below). Thus, a major aim of CCR is not to have a list of relations to annotate cases encountered in corpus data, but to provide a framework of coherence relations explaining differences and similarities between them in cognitive terms. In other words, cognitive validity is one of the main tenets of this model. For this reason, it has been first and foremost used in psycholinguistic studies to explain the way readers process discourse relations (see Chapter 6), and the order in which children acquire them (see Chapter 8). It has also been used to annotate corpus data, for example, the DiscAn corpus³ (Sanders, Vis & Broeder, 2012).

The four primitives suggested by Sanders, Spooren and Noordman (1992) all correspond to a relational criterion. This criterion emerged from the observation that a discourse relation provides more information than the two related segments in isolation, what the authors call their informational surplus. This supplementary information can be categorized into four dimensions that, when put together, constitute the meaning conveyed by each discourse relation.

The first dimension is called polarity, and it separates positive from negative relations. A positive relation functions between the content of the two related segments. For example, in (14), the link is established

³ <https://dev.clarin.nl/node/4198>.

between the fact of winning the competition and the state of happiness. Typically, such relations are conveyed by connectives like *and* or *because*. A relation is negative if it holds between a negated version of one of the segments, as in (15). In this example, Ann's happiness leads to an expectation that something positive happened to her, but this expectation is denied in the second segment. The dimension of polarity separates adversative, concessive and contrastive relations that are all negative from all the other relations that take a positive polarity.

(14) Ann is happy because she won the competition.

(15) Ann is happy but she lost the competition.

[constructed examples]

The second dimension, called basic operation, separates causal from additive relations. On the one hand, causal relations have an implicational order between the segments, as in (14) where the fact of winning the competition implies the state of happiness. In addition to causal relations, conditional relations also have an implicational order, but the difference between them is the status of the cause (hypothetical or real). On the other hand, additive relations do not have an implicational order, but are simply linked by a logical conjunction, as in (16). In this example, the two facts about Elsa only add up and lead to a same conclusion, for example, that Elsa is a gifted person, but there is no implicational order between them.

(16) Elsa is very good at math and she won a swimming competition.

[constructed example]

In addition to additive relations, temporal relations of sequence (conveyed by connectives like *and then*) or temporal overlaps (conveyed by connectives like *meanwhile*) are also linked by a conjunction rather than an implication. But contrary to additive relations, they are also temporally ordered (see below). The criterion of causal vs. additive link also applies to negative relations. For example, the relation of concession in (15) implies the negation of a causal link, whereas the relation of contrast in (17) does not.

(17) Elsa is very good at math but her sister is not.

[constructed example]

The third dimension corresponds to the source of coherence, and separates the objective from subjective relations (originally called “semantic” and “pragmatic” relations in Sanders, Spooren and Noordman (1992) and later relabeled in the psycholinguistic literature

and in the more recent versions of the model). Objective relations are connected at the level of their propositional content, or in other words they concern real-world events not actively constructed by the speaker, as in (18).

- (18) The door slammed because there was strong wind outside.
[constructed example]

In subjective relations, the speaker is actively involved, as it presents a reasoning or speech act that they perform in one or both segments, as in (19). In this example, the fact that the lights are always out does not cause the neighbors' holiday but merely the speakers' conclusion that they are away. This dimension separates subjective relations such as evidence and justification from objective ones such as temporal sequence or cause-consequence.

- (19) The neighbors must be on holiday, because their lights are always out.
[constructed example]

The last dimension applies only to causal relations. It was originally called basic versus nonbasic relations, but has more recently been renamed implicational order (Sanders et al., 2021). In relations that involve an implicational link between the segments, this link can be conveyed in basic order as in (20) or in nonbasic order as in (21).

- (20) Peter was tired so he went home early.
(21) Peter went home early because he was tired.
[constructed examples]

Relations with a basic order present the information in the text following the order of the implication, for example, with the antecedent in the first segment and the consequent in the second segment, as in (20). The order is reversed in nonbasic relations like (21), where the consequent is presented first and the antecedent second in the text.

A summary of all four dimensions and the corresponding relations is reproduced in Table 2.4.

A fifth dimension was added later on to the model in order to account for temporality; in other words the fact that the two segments are ordered in time or not (Evers-Vermeul, Hoek & Scholman, 2017; Sanders et al., 2021). For some relations like additive relations, there is no such order. These are therefore called nontemporal relations. When there is a temporal order, it can either be chronological when the first event chronologically is presented before the second one as in (22) or

Table 2.4 *Taxonomy of relations from Sanders, Spooren and Noordman (1992: 11)*

Basic operation	Source of coherence	Order	Polarity	Class	Relation
Causal	Semantic	Basic	Positive	1	Cause-consequence
Causal	Semantic	Basic	Negative	2	Contrastive cause-consequence
Causal	Semantic	Nonbasic	Positive	3	Consequence-cause
Causal	Semantic	Nonbasic	Negative	4	Contrastive consequence-cause
Causal	Pragmatic	Basic	Positive	5a 5b 5c	Argument-claim Instrument-goal Condition-consequence
Causal	Pragmatic	Basic	Negative	6	Contrastive argument-claim
Causal	Pragmatic	Nonbasic	Positive	7a 7b 7c	Claim-argument Goal-instrument Consequence-condition
Causal	Pragmatic	Nonbasic	Negative	8	Contrastive claim-argument
Additive	Semantic	–	Positive	9	List
Additive	Semantic	–	Negative	10a 10b	Exception Opposition
Additive	Pragmatic	–	Positive	11	Enumeration
Additive	Pragmatic	–	Negative	12	Concession

anti-chronological when the second event is presented before the first one as in (23).

(22) Sam had his breakfast and then he left for work.

(23) Sam left for work after taking his breakfast.

[constructed examples]

One of the key aspects of this framework from a cognitive perspective is that for each dimension, one of the two possible values is deemed to

be cognitively more complex than the other. For example, constructing a causal relation is a more complex cognitive procedure than merely conjoining segments, as it implies constructing an implicational order, often based on world knowledge. Similarly, inferring a subjective relation is more complex than an objective one, because it requires the ability to infer the mental states of the speaker (Zufferey, 2010), an ability known in cognitive psychology as having a theory of mind. Similarly, having to infer a nonbasic order relation is more complex than a basic-order one, and a nonchronological temporal relation is more complex than a chronological one, because in such cases, the chronological or implicational order of the relation reverses the order of presentation in the text.

As mentioned above, the CCR framework places a lot of weight on cognitive validity. Many studies involving language processing and acquisition have found evidence in favor of this model. First, in the domain of language processing, cognitively simpler relations are processed more quickly compared to more complex ones. For instance, causal relations are processed more quickly compared to concessive relations (Köhne & Demberg, 2013), objective causal relations are processed more quickly than subjective ones (Canestrelli, Mak & Sanders, 2013), and causal relations with a basic order are processed more quickly than those with a nonbasic order (Noordman & de Blijzer, 2000). In the field of language acquisition, children start producing cognitively simpler relations like additive and causal relations before more complex ones such as concessive and adversative relations (Evers-Vermeul & Sanders, 2009), master objective relations before subjective ones (Zufferey, 2010; Evers-Vermeul & Sanders, 2011), and understand temporal relations in chronological order better than in antichronological order (Pyykkönen and Järvikivi, 2012).

The CCR framework has also been applied to annotate a range of different corpus data in different languages like Dutch (Sanders, Vis & Broeder, 2012), Spanish (Santana et al., 2018), and Mandarin Chinese (Li, Evers-Vermeul & Sanders, 2013). It has also been used successfully to help nontrained and nonexpert annotators annotate coherence relations using a stepwise approach corresponding to the different dimensions of this model (Scholman, Evers-Vermeul & Sanders, 2016). Finally, it has also been used to annotate corpora with children's productions (e.g., Van Veen, 2011) and parallel corpora (Hoek et al., 2017).

In a nutshell, CCR provides a radically different framework compared to the ones presented so far, in that it does not provide a list of

discourse relations but a decomposition into a set of primitives with the aim of providing a cognitively grounded model of discourse coherence.

2.6 CAN DIFFERENT FRAMEWORKS COMMUNICATE?

All the frameworks for the annotation of discourse relations that we have discussed in this chapter have been used for the annotation of large portions of corpus data. This represents a valuable source of information for researchers working on discourse-related issues. However, one of the limitations of these resources is that they have been annotated with the different labels used in each model, and it is not clear how these labels can be compared across frameworks. Yet, putting these resources together would represent a major step forward for research on discourse-related issues. For this reason, in recent years, there have been some attempts to provide ways to make frameworks more compatible in the future, and to find ways to translate existing annotations from one framework to another. We discuss them in this section.

As Benamara Zitoune and Taboada (2015: 148) observe, there are several problems involved in the comparison of discourse relations across frameworks, such as differences in segmentation, different labels used for discourse relations, and differences in the type of discourse structures that have been annotated. In this section, we will leave aside issues related to segmentation, and focus on the differences of granularity between different models. Several proposals have been made to circumvent it. On the one hand, minimal lists of core relations that are absolutely necessary to annotate discourse relations, and that are robust across languages and genres, have been proposed. For instance, Benamara Zitoune and Taboada propose a hierarchy with three different levels of granularity inspired from the RST, SDRT and PDTB frameworks, leading to a total of 26 relations. They test the validity of their proposal by mapping an RST corpus and two SDRT corpora with this new taxonomy. They report that most mappings were quite straightforward, but there were also problems of granularity that could not be resolved, for example, when a relation was missing in one of the taxonomies or was too fine-grained to have an exact equivalent in the new taxonomy. Another similar attempt at providing a unified taxonomy was conducted as part of the ISO standard for semantic annotation (Bunt & Rashmi, 2016). This proposal contains a set of 20 core discourse relations that are not ordered into

Table 2.5 *Comparison between frameworks using the CCR dimensions*
(Sanders et al., 2021)

Framework Label	Polarity	Basic operation	Implicational order	Source of coherence	Temporality
PDTB <i>cause result</i>	positive	causal	basic	objective	chrono.
SDRT <i>consequence</i>	positive	causal	basic / non-basic	objective / subjective	chrono. / antichrono.

a hierarchy in order to avoid problems of divergent groupings between the frameworks. In both cases, the effort goes in the direction of simplification, but it is not clear yet whether these new taxonomies can really account for all the cases of discourse relations and genres, as is their objective, and whether they will be used for future large-scale annotation projects. It is indeed doubtful that data will be reannotated with the risk of not answering the initial research question anymore.

Instead of providing a new set of coherence relations, Sanders et al. (2021) suggested using the dimensions from the CCR framework as an interlingua to make other annotation frameworks communicate. They therefore decomposed all relations from the RST, SDRT and PDTB frameworks in terms of the five dimensions listed above (see Section 2.5). Thus, each relation receives values in all dimensions, and this makes them comparable independently of the label used in various frameworks. Table 2.5 illustrates how this comparison applies to the relation of contingency–cause–result from the PDTB and the relation of consequence from SDRT:

Given their labels, both relations might appear to cover similar cases, but the decomposition into dimensions shows that this is not the case, as they differ in three out of five dimensions: implicational order, source of coherence and temporality. It appears that the SDRT relation of consequence is more general than the PDTB result relation, as it also encompasses what in the PDTB would correspond to two additional relations: the relation of cause-reason to account for the nonbasic order of the segments, and the relation of justification to account for subjective causal relations.

Thanks to this decomposition, it becomes immediately clear which aspects of the relations vary between frameworks, something that is rather opaque based on their labeling alone. However, the five dimensions do not contain enough characteristics to specify all the features of some relations such as list or condition. For this reason, a limited set of

additional features has been added to the model for comparison purposes. One such feature is conditionality, a feature that applies to the consequence relation in SDRT but not to the causal relations in PDTB.

Finally, in another more limited comparison between the CCR framework and the PDTB-3 taxonomy, Rehbein, Scholman and Sanders (2016) found that the comparability between them was very good, as the modifications made to PDTB-3 compared to PDTB-2 resolved many problems to compare it to other frameworks. They applied both frameworks to a corpus of spoken data consisting of telephone conversations and broadcast interviews, and found that both could adequately account for many relations found in this mode as well.

To summarize, comparing annotations performed between frameworks remains an important challenge that will need to be addressed in the future so that the many available resources can be reused across projects. In this respect, the decomposition of relations into dimensions that can be compared across frameworks seems to be a promising step forward, because of the possibility it offers of abstracting away from relations' labels and listing their core characteristics.

2.7 SUMMARY

The goal of this chapter was to present the main characteristics of four leading models for discourse annotation. We have seen that a major difference between them lies in their scope. While the Rhetorical Structure Theory (RST) and the Segmented Discourse Representation Theory (SDRT) models aim at providing a full-fledged representation of text structures, the Penn Discourse Treebank (PDTB) framework is lexically grounded and theory neutral. The Cognitive Approach to Coherence Relations model (CCR) takes yet another perspective, as it does not aim at listing all the possible relations, but rather at characterizing them by using a set of basic dimensions that are cognitively motivated. Contrary to the other frameworks, it thus favors cognitive plausibility over descriptive adequacy.

The role played by connectives is also quite different across models. While it is quite peripheral in RST, SDRT and CCR, which aim most of all at representing the coherence created by discourse relations, it is central in the PDTB model. The advantage of models like RST and SDRT is that they provide a very comprehensive view of all the other linguistic means that are used to convey discourse relations in addition to connectives. The main advantage of the PDTB is that it provides

thousands of occurrences of connectives annotated with a sense tag. It is therefore the most comprehensive resource available to date to study this class of lexical items. The main advantage of CCR is its psychological grounding that makes it particularly well suited in psycholinguistic studies.

In short, each model has its own advantages and limitations, but each one has been used to annotate data from various languages and genres, as we have illustrated throughout this chapter. The choice of one model over another therefore depends on the goals of the annotation, and more generally on the research questions addressed in a project. An important step ahead in future years will be to find ways to make data annotated with discourse relations more comparable across corpora, either by agreeing on a standardized set of discourse relations, or by finding ways to make different frameworks communicate, for example, by comparing the different dimensions involved in each relation, as proposed by the CCR framework.

DISCUSSION POINTS

- What are the main differences between the RST and SDRT frameworks?
- What are the main advantages of decomposing relations into different dimensions rather than simply listing them in a taxonomy according to the CCR framework?
- Imagine that you plan to annotate the acquisition of discourse relations in a spoken corpus of bilingual children speaking English and Spanish in order to assess whether the order of acquisition is the same in both languages. Which framework would you choose to perform your annotation and why?

FURTHER READING

The definite reference for Rhetorical Structure Theory is Mann and Thompson (1988). For SDRT, the most complete reference is Asher and Lascarides (2003). An introduction to the PDTB framework can be found in Prasad, Webber and Joshi (2017), and an introduction to the novelties of the PDTB-3 annotation is Prasad et al. (2018). A description of the technical aspects of annotation for all these frameworks and many more can be found in the book edited by Ide and Pustejovsky

(2017). The CCR framework and its usefulness for comparing annotations from various sources is presented in Sanders et al. (2021). The original CCR paper that launched this framework is Sanders, Spooren and Noordman (1992). The issue of discourse segmentation is discussed by Hoek, Evers-Vermeul and Sanders (2018) for written discourse and by Degand and Simon (2009) for spoken discourse.