

Making sense of discourse

On discourse segmentation and the
linguistic marking of coherence relations

Published by
LOT
Trans 10
3512 JK Utrecht
The Netherlands

phone: +31 30 253 6111

e-mail: lot@uu.nl

<http://www.lotschool.nl>

Cover illustration: 'Creating structure' by Jasper Stenger & Jet Hoek

ISBN: 978-94-6093-294-6

NUR 616

Copyright © 2018: Jet Hoek. All rights reserved.

Making sense of discourse
On discourse segmentation and the linguistic
marking of coherence relations

Discourse doorzien
Over de segmentatie en linguïstische
markering van coherentierelaties

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor
aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. H.R.B.M. Kummeling,
in gevolg van het besluit van het college voor promoties
in het openbaar te verdedigen
op vrijdag 21 september 2018
des middags te 4.15 uur

door

Jet Hoek
geboren op 4 juli 1990 te Arnhem

Promotoren: Prof. dr. T.J.M. Sanders
Prof. dr. S. Zufferey

Copromotor: Dr. J. Evers-Vermeul

This project was funded through the Swiss National Science Foundation, Sinergia project MODERN (CRSII2_147653).

Voor opa

Contents

Acknowledgments	xi
1 Introduction	1
1.1 Making sense of discourse.....	1
1.2 Discourse in machine translation.....	4
1.3 Research questions.....	7
1.4 Methodological considerations.....	9
1.5 Reading note.....	12
2 Segmenting discourse: Incorporating interpretation into segmentation?	13
2.1 Introduction.....	14
2.2 The clause as the basis for identifying discourse segments.....	16
2.3 Conceptual dependency and the segmentation of embedded clauses.....	20
2.3.1 Clausal complements.....	21
2.3.2 Clausal subjects, restrictive adverbial clauses, and restrictive relative clauses.....	26
2.4 Stance markers and discourse segmentation.....	29
2.4.1 Complement-taking predicates as stance markers.....	30
2.4.2 Stance adverbials and segmentation.....	32
2.4.3 Complement-taking verbs and discourse segmentation.....	33
2.5 Discussion and conclusion.....	35
3 Annotating discourse: Using the Cognitive approach to Coherence Relations for discourse annotation	39
3.1 Introduction.....	40
3.2 The Cognitive approach to Coherence Relations.....	41
3.2.1 POLARITY.....	43
3.2.2 BASIC OPERATION.....	44
3.2.3 SOURCE OF COHERENCE.....	44
3.2.4 ORDER OF THE SEGMENTS.....	45
3.3 Extensions of the original CCR taxonomy.....	45
3.3.1 Additional distinctions within original primitives.....	46
3.3.1.1 TEMPORALITY.....	46
3.3.1.2 VOLITIONALITY.....	48
3.3.1.3 PURPOSE.....	48
3.3.1.4 DIRECTNESS.....	49
3.3.2 Proposing a new distinction: DISJUNCTION.....	52
3.3.2.1 Disjunctions in CCR.....	52
3.3.2.2 Disjunction as a new distinction in CCR.....	53

3.3.3	Operationalizing SOURCE OF COHERENCE: segment-internal distinctions.....	54
3.3.3.1	Identity of the <i>Subject of Consciousness</i>	55
3.3.3.2	Explicit presence of the <i>Subject of Consciousness</i>	56
3.3.3.3	Propositional attitude of the segments.....	57
3.3.3.4	Determining SOURCE OF COHERENCE.....	57
3.3.4	State-of-the-art CCR for discourse annotation.....	57
3.4	CCR as an annotation scheme in practice.....	61
3.4.1	Calculating inter-annotator agreement when using CCR.....	61
3.4.2	Independence of primitives and the use of end labels in addition to primitives.....	63
3.4.3	SOURCE OF COHERENCE versus truth-value.....	65
3.5	Conclusion.....	67

4 Cognitive complexity and the linguistic marking of coherence relations..... 69

4.1	Introduction.....	70
4.1.1	Expectedness and the marking of coherence relations.....	71
4.1.2	Relational and syntactic dependency and the marking of coherence relations.....	73
4.1.3	Outline.....	75
4.2	Cognitive complexity of coherence relations.....	75
4.2.1	POLARITY.....	76
4.2.2	BASIC OPERATION.....	77
4.2.3	SOURCE OF COHERENCE.....	78
4.2.4	ORDER OF THE SEGMENTS.....	80
4.3	Parallel corpus study.....	80
4.3.1	Implication and implicitness.....	80
4.3.2	Corpus.....	82
4.3.3	Annotation of relational dependency.....	84
4.3.4	Inter-annotator agreement of source text annotation.....	84
4.3.5	Translation spotting.....	86
4.4	Results.....	90
4.5	Discussion and conclusion.....	94

5 How connectives and segment-internal elements interact in the linguistic marking of coherence relations..... 99

5.1	Introduction.....	100
5.1.1	The marking of coherence relations by connectives.....	101
5.1.2	The marking of coherence relations by other linguistic elements.....	102
5.1.3	The interaction between connectives and segment-internal elements.....	104

5.2	The marking of coherence relations in parallel corpora	104
5.2.1	Coherence relations in translation	105
5.2.2	The Cognitive approach to Coherence Relations	106
5.3	Division of labor	107
5.3.1	Division of labor between connectives and segment-internal elements	108
5.3.2	Division of labor in translation	111
5.4	Agreement	114
5.4.1	Agreement between connectives and segment-internal elements	114
5.4.2	Agreement in translation	117
5.5	General collocation	120
5.5.1	General collocation between connectives and segment-internal elements	120
5.5.2	General collocation in translation	121
5.6	Discussion and conclusion	123

6	Restrictive relative clauses as discourse segments?	127
6.1	Introduction	128
6.2	Implicit causality	132
6.3	Continuation experiment 1	133
6.3.1	Participants	135
6.3.2	Materials	135
6.3.3	Procedure	135
6.3.4	Annotation and data clean-up	135
6.3.5	Analysis method	136
6.3.6	Results	136
6.3.7	Discussion	138
6.4	Continuation experiment 2	138
6.4.1	Participants	139
6.4.2	Materials	139
6.4.3	Annotation and data clean-up	139
6.4.4	Results	140
6.4.5	Discussion	142
6.5	Experiment 3: Self-paced reading	143
6.5.1	Participants	144
6.5.2	Materials	144
6.5.3	Procedure	146
6.5.4	Data clean-up and analysis	146
6.5.5	Results	147
6.5.5.1	Reading times	147
6.5.5.2	Reaction times on verification statements	148
6.5.6	Discussion	148

6.6	Experiment 4: Eye-tracking-while-reading	150
6.6.1	Participants	151
6.6.2	Materials	151
6.6.3	Procedure	152
6.6.4	Data clean-up and analysis	153
6.6.5	Results	153
6.6.6	Discussion	155
6.7	General discussion and conclusion	156
7	Conclusion and discussion	159
7.1	Overview of main findings	160
7.1.1	Discourse segmentation	160
7.1.2	The linguistic marking of coherence relations	163
7.2	Discourse in machine translation	167
7.3	Topics for discussion and further research	170
7.3.1	Parallel corpora	170
7.3.2	The relationship between hierarchical discourse structure and relation type	172
7.3.3	Expectations about hierarchical discourse structure	173
7.4	A cognitive theory of coherence relations	173
7.5	Conclusion	176
	Appendix A: Sources of examples	179
	Appendix B: Experimental items experiments 1 & 2, Chapter 6	185
	Appendix C: Experimental items experiment 3, Chapter 6	189
	Appendix D: Experimental items experiment 4, Chapter 6	197
	References	209
	Nederlandse samenvatting	229
	Curriculum Vitae	237

Acknowledgments

It's hard to know what to expect when you start your PhD project, but in most respects, the last few years exceeded any expectations I had. One of the things I had not fully anticipated was that it would be this strange mix of having to motivate myself over a *long* period of time while at the same time working with others so much more than I was used to. Here, I would like to thank the people who helped me navigate this, who supported me in getting this book finished, and who contributed toward making my PhD experience —overall speaking (trying to prevent rosy retrospection here)— so enjoyable.

First of all, my advisors: Ted Sanders, Jacqueline Evers-Vermeul, and Sandrine Zufferey.

Ted, I'm so grateful you gave me the opportunity to pursue a PhD. Thank you for sharing your discourse knowledge with me and for making the right calls at the right times throughout my project. As I've told you before, I am really glad that, mostly because of MODERN and TextLink, we got to travel a lot together. This allowed me to get to know you much better (and to discuss with you so many more complex relations that you wished were fun causal examples) than I would have if I had to rely on our scheduled meetings and the rare Trans-encounter.

Jacqueline, I can honestly say that I don't think this book would be here now if you hadn't agreed to become my daily supervisor, and I can't thank you enough for that. Your enthusiasm and encouragement have been so valuable over these last few years, as well as your almost always being available for a quick question, some trouble-shooting, or just a chat.

Sandrine, my project really began to take shape when you introduced me to parallel corpora and all their promising possibilities. The parallel corpus study that was supposed to be only a small part of my project grew and grew and in the end became the very core of this dissertation. Thank you for taking that journey with me and for annotating and discussing so many difficult Europarl examples with me. Throughout my project, I also benefitted greatly from your pragmatic attitude and composure, both of which I admire a lot.

Many thanks also to all the other people I got to work with over the course of my PhD project.

Hannah Rohde, I have really enjoyed working with you for the relative clause part of my project. Thank you for helping me design the different experiments, for helping me grapple with the sometimes confusing predictions, and for helping me solve all the puzzles we encountered along the way. I so look forward to continue working with you!

xii *Acknowledgments*

Merel Scholman: I'm so glad to have had you along for the ride as the Peppi to my Kokki (or, depending on which criteria you use, the Kokki to my Peppi). Thank you for helping me deal with so many problems, be they coherence-related, statistics-related, inter-personal, or personal. I truly hope we can continue to be in each other's lives; if not as colleagues (I'm not in denial, but not entirely not in denial either), as friends.

Many thanks to all the members of the MODERN project, for all the interesting and fun meetings we had and for teaching me more about machine translation than I ever thought I'd know. Special thanks to our project leader, Andrei Popescu-Belis, for managing these different perspectives in a single project, for always thinking about the bigger picture, and for your highly contagious enthusiasm.

I feel also very grateful toward TextLink (led by Liesbeth Degand) and all its participants. Getting to meet so many researchers working on discourse and discussing DRDs and other discourse-related topics within this network has been invaluable to my PhD project and, more generally, my venture into and the field of discourse.

I would also like to thank Kate Cain and the Psychology Department at Lancaster University for their hospitality in October 2017 and for allowing me to run my eye-tracking experiment as fillers in their experiment. Thanks also to Liam Blything, Gill Francey, and Nicola Currie for their help in testing participants.

I also want to thank my (ex-)colleagues from the Uil-OTS and/or Taalbeheersing, among whom other PhD candidates and post-docs: Renske Bouwer, Monica Koster, Louise Nell, Björn 't Hart, Anne van Leeuwen, Brigitta Keij, Jolien Scholten, Suzanne Bogaerds-Hazenbergh, Yuan Xie, Dominique Blok, Myrthe Bergstra, Martijn van der Klis, Stavroula Alexandropoulou, Shuangshuang Hu, Ileana Grama, Anja Goldschmidt, Maartje Schulpen, Heidi Klockmann, Carolien van den Hazelkamp, Frederike Groothoff, Zenghui Liu, Anne-France Pinget, Naomi Kamoen, Marijn Struiksma, and Hannah de Mulder.

Special thanks also to my fellow discourse PhDs and frequent Codex visitors: Yipu Wei, Andrea Santana C., Nina Sangers, Fang Li, Theresa Redl, Philippine Geelhand de Merxem, and Hongling Xiao. I have enjoyed learning about your research, bouncing my own ideas off of you, and just in general, your company in the wonderful world of discourse!

And of course my office mates for the vast majority of my PhD project: Marloes Herijgers and Suzanne Kleijn. Thanks for dealing with me and my messy desk-ways (which was mostly hard on Suzanne), and for all the coffee, tea, and cookies over the years. I guess I'll forgive you for the far insufficient number of ice cream runs and Friday afternoon drinks (but I still think we should keep working on that).

Finally, I would like to thank the people outside of the university who provided me with the necessary support and distraction.

To all my friends outside of the university, and to Anne, Jasper, Andrew, and Renée in particular: thank you for having drinks with me, for barbecuing with me, for taking trips with me, and just in general, for putting up with me. This is usually the place where people promise that they will now have so much more time to spend with their friends, but since we'll soon face several new geographical issues I'll just say that I'm glad to know that even if we don't see each other as much as we'd want to, we're good.

My family, and especially papa, mama, Willem, opa, and oma. Thank you for providing me with a solid home base that I can always fall back on. Also, for supporting me in everything I do while at the same time never being enthusiastic to the point that it becomes downright unbelievable. I wouldn't have it any other way.

And finally, Jelmer. Thank you for having more faith in me than I could ever have in myself, for being at least semi-successful in teaching me to cut myself some slack every once in a while, for finding my rants mostly amusing, for almost always managing to cheer me up, and for so, so much more. You make everything better.

1 Introduction

1.1 Making sense of discourse

When people read or listen to a discourse, they – generally speaking – try to make sense of it. To fully understand a discourse, it is essential not only to know the meaning of each individual clause, but also to figure out how all clauses are related to each other. If all goes well, language users end up with an accurate representation of the discourse. An important aspect of building a mental representation of a discourse is inferring *coherence relations* between *discourse segments* (e.g., Hobbs 1979, Kehler 2002, Sanders, Spooren, & Noordman 1992). For instance, to comprehend the fragment after the italicized context in (1) – and to successfully make the ice cream – it is crucial to infer that the first three discourse segments in the sentence are ordered in time and thus should be executed in that specific order. The final segment of the sentence expresses the result of the first three segments combined.¹

- (1) *All you need to make this ice cream is a few cups of heavy cream and a can of sweetened condensed milk.*
[Whip the cream,] [fold it into the sweetened condensed milk,] [freeze for a few hours,] and [sweet ice cream bliss is yours.]

When inferring a coherence relation, language users have to deduce whether two or more chunks of text constitute, for instance, a cause-consequence relation, a rule and an exception, alternatives, etc. A coherence relation can be explicitly marked by a *connective* (e.g., *because, although*) or a *cue phrase* (e.g., *except for the fact that, by comparison*), as in (2), but this need not be the case. In many instances, a coherence

¹ Sources for all examples can be found in Appendix A.

2 Chapter 1

relation has to be established without the instructions connectives provide – as in (1), where there is no connective between the first and second, and second and third segment – or with limited instructions; the connective between the third and fourth segment in (1), *and*, encodes a relation of addition, but the reader infers a more specific relation of *result*.

- (2) **If** [your freezer happens to be full of garlic and onions,] [it's not the best environment for your ice cream.]

In addition to determining the type of relation, language users have to identify the parts of the discourse between which the relation holds. In (1), the *result* relation links the segment following *and* to the three previous discourse segments combined, rather than to, for instance, only the preceding segment. In research on discourse coherence, the processes of determining the type of relation and identifying the relevant discourse segments are captured, respectively, by the practices of *discourse annotation* and *discourse segmentation*. Segmenting a text and annotating the relations between the segments results in a depiction of a discourse structure; which can provide insight into how people build a discourse representation. A lot of attention has been paid to the types of relations people infer (e.g., Asher & Lascarides 2005, Carlson & Marcu 2001, Hobbs 1990, Kehler 2002, PDTB Research Group 2007, Reese, Hunter, Asher, Denis, & Baldrige 2007, Sanders et al. 1992, Wolf & Gibson 2005), with discourse annotation frameworks that differ in the types and number of relations they distinguish, the granularity of the relation labels, and the criteria relations have to meet in order to be included in a framework's relation inventory. How to segment a discourse, on the other hand, has been studied much less extensively, even though there seem to be several issues with existing segmentation guidelines (formulated by for instance Carlson & Marcu 2001, Mann & Thompson 1988, Reese et al. 2007, Sanders & van Wijk 1996, Wolf & Gibson 2005).

Consider (3), which consists of the first few lines from a blog post.

- (3) I live in what is, by all reasonable standards, a very small apartment. I say this not to complain -- I like my apartment! -- but only to suggest that I am not necessarily in a position to devote an entire cabinet to an ice cream maker that I have never used.

But there it sits, a full-sized ice cream maker, in a cabinet above my three-quarter sized fridge, where it has rested, untouched by milk or sugar, for approximately six years.

Objectively, it makes no sense to keep the ice cream maker, given that I have extremely limited space, and also have never made ice cream. My boyfriend brought the ice cream maker into the relationship. He has never made ice cream, either.

The text in (3) consists of several sentences. It contains a few connectives and cue phrases (e.g., *but*, *given that*), but certainly not all coherence relations are accompanied by an explicit linguistic cue. In fact, it seems that most of the coherence relations in this text have to be inferred without the explicit instructions connectives provide. In addition, establishing the number of discourse segments that make up the text in (3) is not entirely obvious and, as a consequence, the number of coherence relations that hold within the discourse. Most segmentation guidelines take the grammatical clause as the basis for identifying discourse segments, which in (3) would result in fifteen segments. However, all segmentation guidelines have formulated exceptions to the basic clause-as-segment rule. When considering these exceptions, counting the number of discourse segments in (3) becomes much more complicated, especially when it comes to the second sentence. While some approaches consider attribution constructions, which consist of an attribution verb and a clausal complement, such as *suggest that + complement*, to be two separate segments, other guidelines consider these constructions to be a single segment. In addition, almost all frameworks exclude restrictive relative clauses from being discourse segments. In (3), however, it seems plausible to infer that the restrictive relative clause *that I have never used* provides an additional reason for why the author should not have an ice cream maker. This is underlined by the first sentence of the third paragraph, which explicitly states that both the lack of space and the fact that the author has never once made ice cream are reasons for why it makes no sense that she is keeping the ice cream maker. The relative clause can also be manipulated to motivate why it would be perfectly acceptable to keep it despite a lack of space: *I have an extremely tiny kitchen, but it is easy to justify devoting an entire cabinet to an ice cream maker that I use four times a week.*

The fragment in (3), despite being from a seemingly simple text about whether or not to throw away a kitchen appliance, illustrates several issues that will be addressed in the current dissertation: which parts of the text do language users identify and treat as discourse segments? How do they link the correct segments to each other? And how can they infer an appropriate coherence relation between two segments in the absence of a linguistic device that gives them instructions on how to do so? In addition, the fragment demonstrates that connectives mark only a limited number of coherence relations within a discourse, which yields the question of when and why speakers or writers mark a relation by a connective, and when they leave it out.

While the research reported in this dissertation mainly focuses on discourse coherence as a feature of human language, its findings could be of interest to the fields of natural language processing and machine translation (MT). The research project reported in this dissertation was part of a larger project focused on improving MT at the discourse level, both through the use of discourse-level features in MT systems, and by modeling the linguistic devices crucial to discourse coherence (MODERN;

4 Chapter 1

Modeling discourse entities and relations for coherent machine translation²). The studies reported in this dissertation contribute to the latter goal of the project. Furthering our understanding of discourse phenomena in actual language use can function as a basis for improving computer-generated translations, since it helps give a more complete overview of the benchmark – human-generated language and translations – MT systems have to live up to, and helps formulate concrete suggestions for improvement. The main findings of the individual chapters will be related to the area of Machine Translation in the concluding chapter of this dissertation (Chapter 7).

1.2 Discourse in machine translation

Discourse coherence and discourse structure are not just important linguistic features of monolingual language use. In translation, the coherence relations and discourse structure of a source text have to be transferred to the target language in order to preserve the overall coherence of the discourse. While this may at times be a challenge for human translators, it appears to be highly problematic for machine translation systems. Most MT systems to date only operate within sentence boundaries and, as such, deal with sentences in isolation (e.g., Bahdanau, Cho, & Bengio 2014, Koehn, Och, & Marcu 2003). Since the larger discourse context is not taken into account, difficulties in the automatic translation of discourse-level phenomena, which often exceed sentence boundaries, are to be expected. Examples of mistakes made by phrase-based statistical MT in the translation of discourse coherence include wrong connective choice, wrong position of the connective, and unjustified implicature, as will be illustrated in this section.³

The example in (4) illustrates a fairly basic discourse-level mistake.⁴ The MT translation uses *weliswaar*, rather than the admittedly more obvious *hoewel* ‘although’ or *maar* ‘but,’ as a translation for *although*. In Dutch, however, *weliswaar* can only be used in combination with *maar* ‘but’ (as a ‘compound connective’); the absence of *maar* is marked by the \emptyset symbol in the example. The mistake in the MT version could be due to the distance between *weliswaar* and *maar* being too large for the MT system to learn their correlation. Chapter 5 of this dissertation explores how certain combinations of connectives and other linguistic items can together signal a coherence relation. This knowledge could be used to help MT systems avoid the mistake in (4) and deal with similar, potentially more complex, translation options.⁵

² SNSF Sinergia Project CRSII2_147653, led by Prof. Andrei Popescu-Belis.

³ The MT examples used in this chapter were generated by Ngoc Quang Luong (Idiap Research Institute). The Moses phrase-based statistical MT system was trained on approximately 1.2 million EN-NL sentence pairs taken from the Europarl corpus (Koehn 2005).

⁴ This chapter only discusses discourse-level errors in the MT examples; this is not to say that all MT examples are devoid of other types of errors, such as word choice, grammatical errors, etc.

⁵ All examples with an ep-number (ep-year-month-day) were taken from the Europarl corpus (Koehn 2005).

- (4) EN **Although** [the European Union has been very active in the nuclear sector,] [up to now there has been a clear failure of the European Union and G-7 policy.] {ep-99-03-10}
- NL **Alhoewel** de Europese Unie zeer actief is in de nucleaire sector heeft het beleid van de Europese Unie en de G-7 tot nog toe volledig gefaald.
- MT De Europese Unie heeft **weliswaar** zeer actief geweest in de nucleaire sector \emptyset tot nu toe is er sprake van een duidelijke falen van de Europese Unie en de G-7 beleid.
'The European Union has WELISWAAR been very active in the nuclear sector, \emptyset up to now there has been a clear failure of the European Union and G-7 policy.'

The example in (5) features a similarly basic error. In this translation, the MT system has used *maar* 'but' to translate *although*. While this is certainly possible, the connective has to be placed *between* the discourse segments, rather than at the head of the relation. Since *maar* in the MT translation in (5) occurs at the beginning of a sentence, the relation is inferred to hold between the segment following *maar* and the discourse preceding the connective, instead of between the two segments given in (5). *Although*, being a subordinating conjunction, can occur either at the head of the relation, or between the two segments. In learning that *maar* can be used to translate *although*, the MT system has not acquired knowledge about the position of the connectives relative to the discourse segments, a crucial feature in establishing an appropriate discourse structure. Chapter 2 explores how discourse segments can be identified and how the appropriate discourse segments can be linked to each other in the discourse structure; its findings could be used to improve the recognition of discourse segments by MT systems.

- (5) EN **Although** [it would have been better if we had voted at lunchtime,] [I am happy to go ahead now.] {ep-00-03-16}
- NL **Ofschoon** we deze stemming beter hadden gehouden rond de middagpauze, vind ik het goed als we het nu doen.
- MT **Maar** het zou beter zijn geweest als we hadden gestemd tegen de middag, ben ik blij om verder te gaan.
'But it would have been better if we had voted at lunchtime, I am happy to go ahead now.'

A more complex issue can be found in (6). The human translation uses *omdat* to translate *because*; the MT version uses *want*. While the MT system is not required to make the same choice as a human translator in each individual instance, in this case the use of *want* is inappropriate. Both *want* and *omdat* are common Dutch causal connectives, and both are frequently used as translations for *because*. The main

6 Chapter 1

difference, however, is that *want* is a connective typically used in subjective contexts, such as claim-argument relations; *omdat* is more commonly used to mark objective causal relations, such as consequence-cause relations (e.g., Sanders & Spooren 2015, Sanders et al. 1992).⁶ In (6), it becomes clear from the context that it is a known fact that the Council has changed its approach. Both the original English *because* construction and the Dutch human translation allow for this interpretation. The MT translation, however, uses *want*, which gives the impression that the speaker is claiming that the approach has changed. As will be touched upon in Chapter 5, there may be clues inside the discourse segments that point toward a causal relation being subjective or objective, which could potentially be exploited to improve connective choice in MT, thus expanding the initial methods proposed by Meyer (2015).

- (6) *Here I would take issue slightly with my good friend Commissioner Liikanen. In his intervention he spoke about the Council changing its approach and dealing with five Member States as if somehow that was an act of planning. It is not. It is a manifestation of failure. [...]*
EN [The approach has changed] **because** [it has failed.] {ep-98-03-31}
NL De benadering is gewijzigd **omdat** zij mislukt was.
MT De aanpak is veranderd, **want** het is er niet in geslaagd.
'The approach has changed, since it did not succeed.'

The problem with the MT translation in (7) is similarly complex. As is indicated by the brackets in the English original, the causal relation signaled by *because* is embedded under *can you imagine*. The human Dutch translation uses *omdat*, a connective that can be embedded under a syntactic construction. As such, the human translation accurately conveys the meaning and discourse structure of the English fragment. The automatic translation, on the other hand, uses the coordinating conjunction *want*, which cannot be embedded. This has crucial consequences for the discourse structure and interpretation of the sentence. In the MT version, *kunt u zich voorstellen* 'can you imagine' has become part of the first discourse segment, as illustrated by the brackets. Creating an accurate representation of the discourse structure, including determining whether or not a relation is embedded, can at times heavily rely on the meaning of a text fragment, as will be elaborated on in Chapter 2 of this dissertation. Since an MT system is unable to interpret a text, choosing an appropriate connective for a certain discourse structure can be problematic.

⁶ Another major difference between *want* and *omdat* is the fact that *want* is a coordinating conjunction and *omdat* a subordinating conjunction. This difference, however, is less relevant to the example in (5) than the objective/subjective difference.

- (7) EN Can you imagine [all automotive plants being given the possibility of compensation] **because** [they are located in the peripheral regions in the European Union]? {ep-97-05-13}
- NL Kunt u zich voorstellen dat [alle autofabrieken de mogelijkheid krijgen tot compensatie] **omdat** [ze gelegen zijn in perifere gebieden van de Europese Unie]?
- MT [Kunt u zich voorstellen dat alle installaties in de automobieliindustrie krijgen de mogelijkheid van compensatie,] **want** [ze zijn gevestigd in de perifere regio's in de Europese Unie?]
'Can you imagine all automotive plants being given the possibility of compensation, since they are located in the peripheral regions in the European Union?'

The fragment in (8) features a conditional coherence relation. This relation is explicitly signaled by *if*. The human translation also uses a connective, but the relation has become implicit in the MT version, as indicated by the \emptyset symbol. Changes in the marking of coherence relations, be it the addition or removal of a connective, are fairly common in translation. However, the implicitation in the MT version of (8) is problematic in the sense that the conditional relation becomes extremely difficult, if not impossible to recover. As such, the meaning of the fragment is drastically altered. While leaving out a connective may seem like a relatively straightforward mistake, it is actually not entirely clear when a relation should be explicitly marked and when a connective is not necessary for language users to be able to infer the appropriate coherence relation. This question is the focus of Chapter 4. Improving our knowledge about the marking of coherence relations in monolingual language use can help determine when a translation should use a connective and when it can leave the relation implicit.

- (8) EN **If** [this were a debate about the press] [we would all have very strong reservations about a public sector press.] {ep-96-09-17}
- NL **Als** dit debat over de pers zou gaan zouden we allemaal ernstige bezwaren hebben tegen een publieke pers.
- MT \emptyset Dat is een debat over de persvrijheid zouden we allemaal hebben zeer sterke bedenkingen over de publieke sector in de pers.
'That is a debate about the press freedom, we would all have very strong reservations about a public sector press.'

1.3 Research questions

The examples in the previous section do not only illustrate that MT systems have trouble dealing with coherence relations, but also that discourse coherence is a

8 *Chapter 1*

complex and multi-faceted linguistic phenomenon. Segmenting and annotating a text is not an easy task, and researchers have to tackle various problems when trying to accurately depict discourse, both when setting up segmentation and annotation guidelines and creating relation inventories to use when annotating, as mentioned in Section 1.1, and in actual segmentation and annotation practice; it is for instance not always straightforward to establish which relation holds between two specific discourse segments or to determine which parts of the discourse are related to each other by a specific connective.

The overarching research questions of this dissertation can be formulated as follows:

- RQ1: Between which parts of a text do people establish coherence relations?
- RQ2: When and why are coherence relations explicitly marked by a connective or a cue phrase?

Chapter 2 explores discourse segmentation and discourse structure, starting from the assumption that implementing segmentation rules should result in text segments that correspond to the units of thought related to each other in the mental representation of a discourse. Using corpus examples, it argues against two segmentation guidelines that were originally proposed by Mann and Thompson (1988) and that have been implemented in many other discourse annotation approaches as well: the treatment of segmentation and annotation as a two-step process, which prevents the circularity of a process in which annotation and segmentation are intertwined (Taboada & Mann 2006), and the completeness constraint, which poses that the segmentation of a text has to include all elements of that text. The chapter demonstrates that accurate segmentation is in part dependent on the propositional content of text fragments, and that completely separating segmentation and annotation does not always yield text segments that correspond to the text units between which a relation holds in the mental representation of a text. In addition, it argues that elements belonging to the propositional content of the discourse should necessarily be included in the segmentation, but that inclusion of other text elements, for instance stance markers, should be optional.

Chapter 3 focuses on discourse annotation. Specifically, it reflects on using the Cognitive approach to Coherence Relations (CCR; Sanders et al. 1992) for discourse annotation. Unlike most other approaches to discourse, CCR does not make use of a set of relation definitions or ‘end labels,’ but defines basic cognitive primitives that can be used to depict different types of coherence relations. The original CCR proposal defined four primitives that are applicable to all or almost all relations, but several additional distinctions have been proposed over the years, most of which are relevant to only a subset of all coherence relations. After giving a state-of-the-art overview of CCR, this chapter will advocate including an additional distinction,

DISJUNCTION, as an additional primitive. Finally, the chapter reflects on the functionality and practical implications of using the full CCR taxonomy as a tool for annotating discourse.

The segmentation and annotation practices outlined in Chapters 2 and 3 are used to annotate a set of English source text (ST) relations from a parallel corpus. This annotated corpus forms the basis for Chapters 4 and 5, both of which look into the marking of coherence relations. Chapter 4 investigates whether the marking of coherence relations is influenced by cognitive complexity, hypothesizing that cognitively simple relations are left implicit more often than relations that are cognitively more complex. It also examines whether the marking of coherence relations is influenced by the position of the relation in the discourse structure.

Chapter 5 explores the marking of coherence relations by linguistic means other than connectives. The chapter aims to develop a systematic way of categorizing segment-internal elements as signals of coherence relations. On the basis of the different ways in which elements inside discourse segments interact with connectives in the marking of coherence relations, a three-way distinction is proposed between *division of labor*, *agreement*, and *general collocation*. Segment-internal elements can function as signals for coherence relations in all three types of interaction, but the mechanism behind it is slightly different for each type. The presence of a segment-internal signal can eliminate the need for, or reduce the likelihood of, the relation being marked by a connective.

Finally, Chapter 6 investigates whether restrictive relative clauses (RCs) can enter into a coherence relation with their host clauses, and if they should be considered to be potential discourse segments. While restrictive RCs are usually excluded from receiving discourse segment status, as is elaborated on in Chapter 2, the parallel corpus study that serves as the basis for Chapters 4 and 5 shows that restrictive RC constructions are used as translations for coherence relations, and vice versa, which suggests that restrictive RCs are, at least sometimes, treated as discourse segments. In four experimental studies, this chapter tests whether restrictive RCs can influence the language users' expectations about the upcoming discourse; specifically, whether expectations differ depending on the type of coherence relation that could be inferred between the restrictive RC and its matrix clause.

1.4 Methodological considerations

This dissertation deals with discourse segmentation, discourse annotation, and discourse structure from both a text-linguistic and a language user perspective, and uses a combination of theoretical exploration, qualitative and quantitative corpus-based methods, and experimental methods to address the individual research questions. In line with Sanders et al. (1992), coherence relations are considered to be cognitive constructs; they are a feature of the cognitive representation of a text, rather

than a feature of its linguistic realization (see also Hobbs 1979, Kehler 2002, Kehler, Kertz, Rohde, & Elman 2008). From this cognitive perspective, the way in which language users process and produce discourse should inform how we treat texts in our linguistic analysis of discourse, and vice versa; linguistic theories about discourse should be cognitively plausible and make reasonable predictions about actual language use.

Chapters 2 and 3 make use of real language excerpts in developing and assessing both theory on discourse coherence and discourse segmentation and annotation practices. In turn, some of the conclusions about discourse analysis practices in these chapters comment on how language users process discourse. For instance, the observation that segmentation and annotation cannot always be entirely separated because the interpretation of the fragment is needed to resolve discourse structural ambiguities (Chapter 2) also bears on the processing of discourse by language users; it predicts that language users, at least sometimes, use their interpretation of a text to connect the right segments to each other. The use of corpus data in Chapters 2 and 3 is monolingual and predominantly qualitative.

Chapters 4 and 5 make use of parallel corpus data. The corpus consists of English ST fragments and translations into Dutch, German, French, and Spanish. Both chapters primarily use translations as a method to research a monolingual linguistic phenomenon from a cross-linguistic perspective, and the use of multiple target languages allows for a cross-linguistic comparison. Both chapters depart from the point of view that if a phenomenon (in this case, the linguistic marking of coherence relations) is governed by general cognitive principles, it should not differ between languages.

Another benefit of using a parallel corpus over a monolingual corpus is that its analyses and results also comment on translation; how do translators deal with a specific linguistic phenomenon, which words or constructions are used as target text (TT) equivalents for a specific ST word or construction, how often do translators omit parts of the ST in the TT, etc.? Gaining a better understanding of how language users establish coherence in monolingual discourse, as well as improving our knowledge about how, and to what extent, discourse coherence is preserved in translation, could be valuable for improving the quality of computer-generated translations.

A final important benefit of using a parallel corpus over a monolingual corpus is that translations can make observable what is left up to the interpretation of the researcher in monolingual texts. As mentioned above, coherence relations are a feature of the cognitive representation of a discourse, rather than of the linguistic realization of a text. That is why even using a monolingual corpus, studying coherence relations requires the annotator's interpretation of the text. A lot of research on discourse coherence therefore involves corpus annotation, in which the coherence relations that hold between the segments in a text are labeled (see also Chapter 3). To test and demonstrate the reliability and reproducibility of the annotation effort,

annotations from one coder can be compared to those from other coders. Discourse annotation is a fairly difficult task anyway, with very high inter-annotator agreement scores being a rarity (e.g., Spooren & Degand 2010), but annotating implicit coherence relations (i.e., relations without a connective) is particularly tricky, and agreement between coders tends to drop considerably for implicit relations (Miltakaki, Prasad, Joshi, & Webber 2004, Prasad et al. 2008). In the parallel corpus study, all annotations are performed on a dataset of explicit English ST relations. Annotation of the relations is therefore more reliable than it would have been if we had directly annotated implicit coherence relations. The marking of coherence relations is then studied using the translations, as is explained in the next paragraph.

The challenge that coherence relations pose for corpus research seems to be part of the more general problem that “the meaning of linguistic expressions is the least tangible of linguistic phenomena” (Geeraerts & Cuyckens 2010:18) and that “meaning is not directly observable” (Noël 2003:758). An alternative method that has been proposed for researching meaning is to make use of parallel corpora (e.g., Dyvik 1998, Melamed 2001, Noël 2003, Teubert 1999). In this approach, the translator is treated as a naive ‘annotator,’ whose main purpose was to accurately convey the meaning of the ST in the TT. Variety in the linguistic means used to arrive at a similar meaning can inform researchers about the meaning of individual elements or constructions. In (9), for example, the original English fragment expresses a negative conditional relation, signaled by *unless*. Out of all target languages, only Spanish uses the same construction to signal the same relation: *a no ser que* is equivalent to *unless*. The other three languages, despite the fact that they have a connective equivalent to *unless*, express the same meaning using different linguistic elements. French uses a combination of *si* ‘if’ and negation, German uses *solange* ‘as long as’ plus negation, and Dutch uses a preposition, *zonder* ‘without’ and expresses the meaning of the first segment in a nominalization (*een nieuwe duurzame landbouwpraktijk* ‘a new, sustainable agricultural practice’).

- (9) EN **Unless** [new, sustainable agricultural outputs are devised,] [farming and the countryside’s problems can only get worse.] {ep-00-10-25}
- NL **Zonder** een nieuwe, duurzame landbouwpraktijk worden de problemen van de boeren en het platteland alleen maar groter.
- DE **Solange keine** neuen, dauerhaft umweltgerechten Produkte entwickelt werden, wird es mit der Landwirtschaft und dem ländlichen Raum nur weiter bergab gehen.
- FR Les problèmes de l’agriculture et du monde rural ne pourront qu’empirer **si** l’on n’élabore **pas** des productions agricoles nouvelles et durables.
- ES **A no ser que** se diseñen nuevas producciones agrícolas sostenibles, los problemas de la agricultura y del campo no pueden sino empeorar.

Despite the use of different constructions, all translations in (9) preserve the meaning of the ST fragment. Using parallel corpora can thus provide insight into the meaning and function of linguistic elements or constructions. The potential of using parallel corpora to research discourse coherence phenomena has already been demonstrated in several other studies (e.g., Cartoni, Zufferey, & Meyer 2013, Cartoni, Zufferey, Meyer, & Popescu-Belis 2011, Hansen-Schirra, Neumann, & Steiner 2007, Levshina & Degand 2017).

Finally, Chapter 6 takes an experimental approach to the question of whether restrictive RCs can have a function at the discourse level. Departing from the assumption that if restrictive RCs can function as discourse segments they should be able to influence discourse-level phenomena, the chapter uses a combination of experimental methods to investigate the influence of restrictive RCs on several different discourse-level phenomena. Two continuation tasks focus on, respectively, expectations regarding which discourse participant will be the focus of the continuing discourse (next-mention bias) and expectations about discourse structure. A self-paced reading task investigates whether restrictive RCs are similar to other types of discourse segments in that they are read slower or faster depending on the type of coherence relation they feature in. Finally, an eye-tracking study examines whether restrictive RCs can influence expectations about upcoming types of discourse relations and, in addition, investigates the influence of restrictive RCs on next-mention biases in online processing. The combination of methods provides converging evidence that suggests that restrictive RCs can indeed have a function at the discourse level and should not categorically be excluded from receiving discourse segment status in discourse segmentation and annotation practices.

A final methodological note concerns spoken versus written language. The specific parallel corpus that is used in this dissertation consists of the proceedings of the European Parliament. All ST fragments are transcriptions of spoken language without accompanying audio files; all translations are written. All experiments in Chapter 6 also feature written language. It should be noted that the discourse phenomena addressed in the different chapters are not expected to differ qualitatively between written and spoken language, and that the results and observations from each chapter are expected to apply to both spoken and written language.

1.5 Reading note

All chapters were written as individual papers. As a result, there is some overlap in theoretical background and method sections. The main benefit of this arrangement is that all chapters can be read in isolation. Publishing information is provided at the beginning of each chapter; while all papers have co-authors, the development of the main ideas, the set-up and execution of the studies, the analysis and interpretation of the results, as well as the writing of the papers have been predominantly my own.

2 Segmenting discourse

Incorporating interpretation into segmentation?

Discourse segmentation is an important step in the process of annotating coherence relations. Ideally, implementing segmentation rules results in text segments that correspond to the units of thought related to each other. This chapter demonstrates that accurate segmentation is in part dependent on the propositional content of text fragments, and that completely separating segmentation and annotation does not always yield text segments that correspond to the text units between which a conceptual relationship holds. In addition, it argues that elements belonging to the propositional content of the discourse should necessarily be included in the segmentation, but that inclusion of other text elements, for instance stance markers, should be optional.

This chapter has been published as:

Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2017). Segmenting discourse: Incorporating annotation into segmentation? *Corpus Linguistics and Linguistic Theory*. Online preview.

2.1 Introduction

Annotated corpora have become increasingly valuable resources for the study of language. They allow us to investigate the functions of linguistic forms, to study the linguistic realization of particular functions, to test linguistic theories, and to develop new ones. Many annotated corpora contain annotations at the levels of syntax, semantics, and morphology, as well as the annotation of lexical features. In addition, the last two decennia have seen the rise of corpora annotated at the level of discourse. At the discourse level, one of the things that are annotated is the coherence within a text. By annotating the *coherence relations* within a discourse, it becomes apparent how idea units in a text are related to each other, e.g., are they causally related, contrasted, part of an enumeration, etc.?

A coherence relation can be defined as “an aspect of meaning of two or more discourse segments that cannot be described in terms of the meaning of the segments in isolation,” or, in other words, the meaning of a coherence relation is “more than the sum of its parts” (Sanders, Spooren, & Noordman 1992:2). In line with Sanders et al. (1992), we consider coherence relations to be a feature of the cognitive representation of a text, rather than a feature of its linguistic realization (see also Hobbs 1979, Kehler 2002, Kehler, Kertz, Rohde, & Elman 2008). This definition assumes coherence relations to hold between the idea units that readers or listeners construct on the basis of the linguistic input. If we want to make claims about the nature of such coherence relations, it seems important that the text segments indicated to feature in a coherence relation correspond to the idea units that are related to each other in the cognitive representation of a discourse.

The notion of *idea unit* is not a clearly delineated linguistic category, unlike for instance the notions of *subject* and *object* in syntactic annotation. While a lot of attention is paid to creating relation inventories specifying the types of relations that can hold between segments (e.g., Asher & Lascarides 2005, Carlson & Marcu 2001, Hobbs 1990, Kehler 2002, PDTB Research Group 2007, Reese, Hunter, Asher, Denis, & Baldridge 2007, Sanders et al. 1992, Wolf & Gibson 2005), much less theoretical consideration has been given to the exact characteristics of the segments and the way in which they are structured in a discourse (notable exceptions are Matthiessen & Thompson 1988, Polanyi 1988, Schilperoord & Verhagen 1998, Verhagen 2001). Many approaches to discourse annotation have taken the clause as the basis for identifying segments, although annotation frameworks are not uniform in this respect and exceptions or addenda to the clause as basic unit also differ between approaches. The variability between annotation approaches in their operationalization of *idea units* and the (syntactic) rules on the basis of which they identify discourse segments has consequences for the eventual annotation of the coherence relations that are annotated in a corpus, and can consequently affect theories and conclusions formulated on the basis of the data.

Taking a syntactic structure as the basis for segmentation rules makes the segmentation process relatively objective and enables annotators to treat segmentation and annotation as separate steps. However, it appears that strict application of these segmentation rules does not always result in segmentation that does justice to the interpretation of a fragment. As will be illustrated in this chapter, applying conventional segmentation rules may produce segments that are too small, in which case it does not include an entire unit of thought, or too big, in which case only part of the segment connects to the adjacent segment. Alternatively, certain inferred coherence relations may not be segmented at all, as is often the case for coherence relations that are embedded in syntactic constructions such as complement clauses or restrictive relative clauses.

The current chapter will theoretically approach discourse segmentation and focus on two issues concerning segmentation that were proposed by Mann and Thompson (1988) in their introduction of Rhetorical Structure Theory, but that have been implemented in many other discourse annotation approaches as well: the treatment of segmentation and annotation as a two-step process, which prevents the circularity of a process in which annotation and segmentation are intertwined (Taboada & Mann 2006), and the completeness constraint, which poses that the segmentation of a text has to include all elements of that text. In this chapter, we consider segmentation to be accurate when the segments correspond to the idea units that are related to each other. We propose that accurate segmentation is at least in part dependent on the propositional content of text fragments, and that completely separating segmentation and annotation, as well as adhering to the completeness constraint, can be at the expense of the quality of the segmentation.

After establishing the clause as the syntactic basis for the identification of discourse segments, we discuss fragments, mainly from the Europarl Direct corpus (Cartoni, Zufferey, & Meyer 2013, Koehn 2005; all fragments were originally uttered in English), that present segmentation difficulties. We focus specifically on fragments with complement structures, sentential adverbs, restrictive relative clauses, and stance markers. Building on a proposal for discourse segmentation by Schilperoord and Verhagen (1998), we present an approach to segmentation that results in text segments that correspond to the text units between which a conceptual relation is presumed to hold. As a means of determining whether segments actually represent the units of thought related to each other, we will make use of paraphrases (see for example Sanders 1997) throughout this chapter. Comparing different paraphrases of the same relation can help determine between which idea units a coherence relation holds. The idea units that feature in the best paraphrase should be represented by the text segments.

2.2 The clause as the basis for identifying discourse segments

The smallest unit that can function as a discourse segment is often taken to be the grammatical clause (e.g., Evers-Vermeul 2005; Mann & Thompson 1988; Sanders & van Wijk 1996; Wolf & Gibson 2005), which can be defined as a unit headed by a verb. This rule was introduced by Mann and Thompson (1988) as a theory-neutral approach to the classification of a text into segments. Considering the definition of coherence relations we employ, selecting the clause as the minimal unit for discourse segments seems appropriate, since the clause is the smallest grammatical unit that can function meaningfully in isolation.

Requiring discourse segments to be minimally clauses eliminates prepositional phrases as discourse units: (1a) is considered to be a single discourse segment, even though its meaning is similar to (1b), which consists of two segments between which a causal relation holds.

- (1a) Their fears and uncertainties have been compounded because of their belief that immigrants will pose a threat to future employment.
{ep-01-02-14}
- (1b) [Their fears and uncertainties have been compounded]_{S1} because [they believe that immigrants will pose a threat to future employment.]_{S2}
- (1c) Their fears and uncertainties have been compounded because of their beliefs.

Although sentences with prepositional phrases can be very similar to coherence relations, as in (1a), this is often not the case. It is, for example, not possible to paraphrase (1c) in a way that resembles a coherence relation, since the prepositional phrase contains only a simple noun phrase.

In addition, employing the criterion that discourse segments have to be clauses eliminates the possibility of considering fragments such as (2a), in which a verb, in this case *cause*, signals causality, as coherence relations. Even though (2a) resembles the causal relation in (2b) in meaning, it is only one clause and does therefore not contain a coherence relation.

- (2a) In the year 2000 smuggling of tobacco caused losses of GBP 3.8 million to the British Exchequer. {ep-02-02-05}
- (2b) [In the year 2000 the British Exchequer lost GBP 3.8 million,]_{S1} because [tobacco was smuggled into the country.]_{S2}

One of the advantages of taking the clause as the basis for identifying units and not considering prepositional phrases and the objects of causal verbs to be independent discourse units is that it allows us to systematically distinguish between intra- and

interclausal ways of expressing something, for instance the causality in the above examples (e.g., Degand 1996, Stukker, Sanders, & Verhagen 2008).

In theory, identifying clauses should be fairly straightforward. However, clauses need not be complete, and although it is commonly agreed upon that clauses with ellipted elements can be discourse segments, there is less consensus on when exactly a clause should no longer be considered to be a discourse segment. Two types of approaches for assigning discourse segment status can be identified: defining what can still be considered a clause, or defining what cannot be considered a clause.

Both Sanders and van Wijk (1996) and Carlson and Marcu (2001) provide guidelines for what can still be considered a clause. Sanders and van Wijk (1996: 126), for example, allow only one “major constituent” to be contracted. Carlson and Marcu (2001: 12) allow the subject, auxiliary verb, and adverb of a clause to be ellipted, and even the main verb, provided that “there are strong rhetorical cues marking the discourse structure.” Neither approach would allow the segmentation in (3).

- (3) The virus harms cold-blooded animals. It does not replicate at temperatures above 25° centigrade and [would,]_{S2a} if [present in fish for human consumption,]_{S1} [be inactivated when ingested.]_{S2b}
{ep 00-03-01}

In the first segment (S_1) of the coherence relation in (3), both the subject and the main verb have been left out, without there being any “strong rhetorical cues.”¹ If we were to adhere to the segmentation guidelines provided by Sanders and van Wijk (1996) or Carlson and Marcu (2001), we would not be able to segment the conditional relation in (3). Not segmenting this relation seems overly conservative, since the segmentation in (3) seems very plausible and exactly captures the two segments related by the connective *if*. Not segmenting the conditional relation would lead to a crucial coherence relation missing from the final annotation of the fragment.

Pander Maat (2002:41), on the other hand, proposes that multiple elements can be contracted in a sentence, as long as in addition to a connective there is also another phrase present between the non-contracted elements. Although it is not entirely clear how this guideline applies to (3), Pander Maat’s segmentation rule appears to be primarily aimed at excluding the possibility of segmenting coordinated nouns, which is not the case in (3). Wolf and Gibson (2005) also seem to prioritize excluding coordinated elements, since they state that they do not consider conjoined nouns in a

¹ Carlson and Marcu (2001: 12) do not give a concrete definition of a ‘strong rhetorical cue,’ but do provide the following example (in bold): “Back then, Mr. Pinter was **not only** the angry young playwright, **but also** the first to use silence and sentence fragments and menacing stares, almost to the exclusion of what we previously understood to be theatrical dialog.” It is not clear whether *if* is a strong enough rhetorical cue, since it only marks one of the discourse segments and is not as prominent as *not only ... but also*.

noun phrase or conjoined verbs in a verb phrase to be separate discourse segments. If we were to follow Wolf and Gibson's (2005) guidelines, (3) could be segmented, since there is no coordination within a phrase.

The type of elision that is illustrated in (3) is not exclusive to conditional relations, but can for instance also be found in segments preceded by *although* or *but*.

- (4) Although [no expert,]_{S1} [I would certainly support the calls for all prisoners of conscience to be freed, in Syria and elsewhere.]_{S2}
 {ep-02-06-13}
- (5) ... [parties can choose their own contract law in relation to these particular contracts,]_{S1} but [not their own winding-up proceedings law.]_{S2} {ep-01-01-05}

As in (3), both the subject and the finite verb have been left out of the clauses following the connective in (4) and (5). Strikingly, in all three fragments, the elided verb is a copula verb. The elements following *although* in (4), *but* in (5), and *if* in (3), are therefore all subject complements and, as such, part of the predicate. If we slightly adjust our definition of a clause from “a structure headed by a verb” to “a structure containing a predicate,” we could formulate the tentative segmentation rule that structures can be discourse segments if they contain (at least part of) a predicate.

If we use the presence of a predicate, or parts of a predicate, as the criterion for discourse segment status, we automatically include non-finite clauses as potential discourse segments. Most discourse annotation approaches seem to indeed allow segments of coherence relations to be non-finite: this is explicitly stated in some manuals (e.g., Carlson & Marcu 2001:6-7) or it can be concluded on the basis of provided definitions and examples (e.g., Mann & Thompson 1988, PDTB Research Group 2007, Reese et al. 2007). At the same time, using the presence of a predicate for assigning discourse segment status excludes structures such as prepositional phrases and non-clausal adverbials or modifiers from receiving discourse segment status, which is also in line with most discourse annotation manuals (e.g., Carlson & Marcu 2001, Mann & Thompson 1988, PDTB Research Group 2007, but not Reese et al. 2007:3).

Taking the predicate instead of the verb as the basis for assigning discourse segment status, prevents compound subjects from being segmented, since subjects are not part of the predicate. Segmenting coordinated nouns in subject position seems indeed something to avoid if discourse segments have to correspond to a unit of thought. Fragment (6), for instance, expresses only one unit of thought, even though it contains a compound subject. The segmentation indicated in (6) does therefore not seem appropriate, which is signaled by the hashtag in front of the fragment.

- (6) # [The Commissioner]_{S1} and [Mr Hatzidakis said that regional disparities will become twice as great.]_{S2} {ep-01-01-31}

Using the presence of a predicate as the basis for segmentation would, however, allow objects to be individual discourse segments. This appears to be too liberal, since it would also allow segmentations like the one in (7), despite the fact that the fragment expresses only one unit of thought: one group of people is being thanked for the same thing.

- (7) # [I want to thank the rapporteur,] [the Commissioner] and [other colleagues who are here tonight.] {ep-97-11-18}

It seems therefore necessary to also include a rule resembling Pander Maat's (2002) or Wolf and Gibson's (2005) in order to prevent segmentation of coordinated structures within a single phrase. By adding this, we exclude segmentation of for instance coordinated nouns, as in (7), or coordinated verbs, as in (8). Amending our predicate-based segmentation rule with the rule that coordinated structures within a single phrase cannot be segmented would exclude segmentations like the ones in (7) and (8), but potentially allow the segmentation in (9).

- (8) # [I, therefore, would ask]_{S1} and [request that this House support Amendment No 4.]_{S2} {ep-97-03-11}
- (9) [I want to congratulate Mrs van den Burg for an enormously well done job]_{S1} and [the Commissioner for introducing this directive.]_{S2} {ep-02-11-20}

The fragment in (9) contains two direct objects, but they are, arguably, not coordinated within a single noun phrase. This results in the segmentation indicated in (9). Unlike (7) or (8), (9) appears to contain two separate idea units, which in this case are explicit speech acts: thanking Mrs. van den Burg for her great output on the one hand, and thanking the Commissioner for coming up with the initiative on the other. The segmentation in (9) therefore seems a more appropriate representation of the discourse structure than the segmentations in (7) or (8). Whether or not two elements are conjoined within a single phrase, and, consequently, whether they should be considered to be independent discourse segments, can be left to the judgment of the annotators.

This section outlined the essential structural properties of discourse segments. The next section will establish another criterion a clause has to satisfy in order to have the status of discourse segment: *conceptual dependency*, which entails that if a clause is an integral part of another clause, it cannot be an independent discourse segment. After introducing the concept of conceptual dependency, we will discuss the

consequences the conceptual dependency criterion has for the process of attributing discourse segment status to clauses, and, in turn, for discourse segmentation.

2.3 Conceptual dependency and the segmentation of embedded clauses

Clauses may satisfy all structural criteria outlined in Section 2.2 and still be excluded from having discourse segment status. The general rule that clauses can be discourse segments is often amended by a few exceptions. The clause types listed in (10) are for instance often denied the status of discourse segments:

- (10)
- i **Clausal complements** (*We saw that people wanted to dance*)
Carlson & Marcu 2001, Evers-Vermeul 2005, Mann & Thompson 1988, Sanders & van Wijk 1996
 - ii **Clausal subjects** (*Dancing is my favorite thing to do*)
Carlson & Marcu 2001, Evers-Vermeul 2005, Mann & Thompson 1988, Sanders & van Wijk 1996
 - iii **Restrictive relative clauses** (*Susan likes men who can dance*)
Evers-Vermeul 2005, Mann & Thompson 1988, Reese et al. 2007, Sanders & van Wijk 1996, Schilperoord & Verhagen, 1998, Verhagen 2001
 - iv **Restrictive adverbial clauses** (*I am going to dance until the music stops*)
Evers-Vermeul, 2005, Pander Maat, 2002, Renkema, 2009, Schilperoord & Verhagen, 1998

Although Reese et al. (2007) do not specifically list clause types excluded from receiving discourse segment status, with the exception of restrictive relative clauses (p. 4), they do state that they do not allow segmentation of embedded structures (p. 3). In practice, this means that at least clausal subjects and clausal complements are also not viewed as discourse segments in their annotation method.

Several approaches to discourse annotation include *attribution relations* in their relation inventory (e.g., Carlson & Marcu 2001, Reese et al. 2007, Versley & Gastel 2013, Wolf & Gibson 2005). Attribution relations indicate who is responsible for the information in a fragment (e.g., Pareti 2012), as in (11).

- (11) You also said that the budget should have the same discipline as national budgets. {ep-99-09-14}

Attribution relations inherently assign discourse segment status to clausal complements. In (11), for instance, *you also said that* would be S_1 of the attribution,

while *the budget should have the same discipline as national budgets*, a clausal complement, would be S_2 . In order to be able to consider attribution relations as coherence relations, some annotation approaches have included exceptions to their segmentation rules (or rules for attributing discourse segment status) for fragments that contain communication verbs. Carlson and Marcu (2001:7) for instance state that “normally, clausal complements are not considered to be EDUs [elementary discourse units – discourse segments]. We make exception to this in the case of clausal complements of *attribution verbs*” (original emphasis). However, neither Carlson and Marcu (2001) nor any of the other annotation approaches provide a comprehensive explanation for making exceptions to segmentation rules on the basis of verb semantics. The definition of coherence relations employed in this chapter seems to exclude attribution relations as coherence relations: the meaning of an attribution construction as a whole is not *more* than the sum of its parts, and only one of the two “segments” of attribution relations can function meaningfully in isolation, namely the embedded clause. The importance of segments being able to function meaningfully in isolation for their status as discourse segments will be further elaborated upon in Section 2.3.1, in which we introduce the notion of conceptual dependency to explain why clausal complements and the other clause types listed in (10) are often excluded from being independent discourse segments.

2.3.1 Clausal complements

Schilperoord and Verhagen (1998) introduce the notion of *conceptual dependency* to explain why embedded clauses are often excluded from being independent discourse segments, something they themselves do not strictly agree with:

If a constituent of clause A is conceptually dependent on a clause B, B is an integral part of the conceptualization of A, and therefore not available as a separate discourse segment (cannot enter into a discourse coherence relation with A, or any other part of the discourse). (p. 150)

Matrix clauses that contain a clausal complement or a clausal subject are not complete without the complement or the subject and are therefore not conceptually independent. Noun phrases that are followed by a restrictive relative clause are also, for their conceptualization, dependent on the restrictive relative clauses: without the restrictive relative clause, the concept to which the noun phrase refers is usually underspecified. Since coherence relations are defined to hold between segments that can potentially be independent (Sanders et al. 1992), there can be no coherence relation between clausal complements, clausal subjects, or restrictive relative clauses, and their host clauses. Crucially, this definition of conceptual dependency assumes that it is the main clause that is dependent on the subordinate clause, instead of the other way around. The subordinate clause from (11), for example, could by itself be an independent

22 Chapter 2

discourse unit, as is illustrated by (12). (12) is a full clause, from which no essential elements are missing.

- (12) The budget should have the same discipline as national budgets.

Schilperoord and Verhagen (1998) point out that not treating the clause types listed in (10) as discourse segments can at times be problematic. They provide the Dutch example in (13), in which dashes are used to indicate clause boundaries, to illustrate that not segmenting embedded clauses can result in a segmentation that underestimates the number of discourse segments in a fragment.

- (13) Daarbij komt // dat zijn vrouw ernstig gehandicapt is // en dat hij een gezin heeft te onderhouden.
Thereby comes // that his wife severely disabled is // and that he a family has to take care of.
To this it can be added that his wife is severely disabled and that he has to take care of his family.
(Schilperoord and Verhagen 1998:145)

(13) contains three clauses, but since two of them are coordinated clausal complements and therefore integrated parts of the main clause, applying the clause criterion results in segmenting (13) as one discourse segment. However, Schilperoord and Verhagen (1998) point out that this goes against the intuition that two idea units are contained in the fragment: *his wife is severely disabled* and *he has a family to take care of*. They propose that after the first complement, the main clause has been completed, and is therefore not conceptually dependent on the second complement clause. The second complement clause can then be treated as a separate discourse segment.

- (13a) [Daarbij komt dat zijn vrouw ernstig gehandicapt is]_{s1} en [dat hij een gezin heeft te onderhouden.]_{s2}

Although this seems like an adequate solution for this particular fragment, problems arise when trying to apply this same line of reasoning to relations such as (14).

- (14) *Mr President, I should like to take Commissioner Bolkestein back to the last part-session here when we discussed sales promotion.*
He may remember that // I complimented him // because he had written an article in a journal // complimenting Parliament on // rescuing the internal market. {ep-02-09-25}

(14) contains five clauses, indicated by dashes, but the main clause, *he may remember that*, is conceptually dependent on a complement. If complement clauses are not allowed to be segmented, (14) would be a single discourse segment, since everything is embedded under the matrix structure *He may remember that*, or, in case of the fourth and fifth clause, embedded under the matrix structure and one or two other structures (as a reduced relative clause modifying the NP *an article in a journal*, and as a complement of the prepositional verb *compliment on* within the reduced relative clause, respectively). Following Schilperoord and Verhagen's (1998) reasoning, we arrive at the segmentation in (14a). Considering only the second clause, *I complimented him* as the complement embedded in the main clause suffices to make the main clause a conceptually independent unit. The clause following *because* can then be considered an independent discourse segment, which means it can enter into a coherence relation with other parts of the discourse. In (14), the coherence relation is explicitly signaled by means of *because*, indicating that the third clause *he had written an article in a journal complimenting Parliament on rescuing the internal market* is a reason for the content of the preceding discourse segment. It seems, however, inaccurate to state that the fact that Commissioner Bolkestein once wrote an article is the reason for him remembering that the speaker once complimented him. Instead, it is more plausible that Commissioner Bolkestein's article was the reason for the speaker to compliment him. If the objective behind discourse segmentation is to represent the units of thought that are related to each other, the segmentation in (14a) seems undesirable, while the segmentation in (14b) more accurately captures the discourse structure.

- (14a) [He may remember that I complimented him]_{S1} because [he had written an article in a journal complimenting Parliament on rescuing the internal market.]_{S2}
- (14b) He may remember that [I complimented him]_{S1} because [he had written an article in a journal complimenting Parliament on rescuing the internal market.]_{S2}

Although the segmentation in (14b) may be appealing on the basis of the propositional content of the segments between which the causal relation is indicated to hold, it does leave the main clause of the sentence stranded. We want to propose that even if a complement is segmented as in (14b), the coherence relation as a whole can function as the complement of the main clause, as in (14c). This makes it structurally identical to a simple complement construction such as (15).

- (14c) [He may remember that [I complimented him]_{S1a} because [he had written an article in a journal complimenting Parliament on rescuing the internal market.]_{S1b}]_{S1}

- (15) He may remember that I complimented him.

Schematically, this can be represented as in (16). X can be a single clause or a bigger chunk of text composed of multiple clauses.

- (16) He may remember that X.

The relation in (17) differs from the relation in (14), even though both fragments have identical surface structures. In (17), the second segment in the coherence relation, *Air France cancelled my flight at 2.10 p.m.*, is not a reason for the content of the complement, the fact that the speaker is present in the meeting, but rather an argument supporting the content of the main clause (including its complement), the statement that it is an achievement that she is present.

- (17) *Madam President, it is in itself an achievement that we are having this debate on the new URBAN Community initiative and [it is an achievement that I am here tonight]_{S1} because [Air France cancelled my flight at 2.10 p.m.]_{S2} but I am here!* {ep-00-02-14}

Adapting the segmentation rules to acknowledge that there can be other idea units expressed in a fragment in addition to the main clause will result in a more complete and accurate description of the discourse as a whole, since no information is lost because of the embeddedness of a clause. At the same time, allowing for the possibility of segmenting embedded clauses enables us to distinguish fragments in which a coherence relation holds between two clauses within a complement, as in (14), from fragments in which a clause is related to a main clause that contains a complement, as in (17). This difference is not only relevant to the organization of the discourse structure, but also helps us differentiate between two distinct meanings. It has been proposed that an important function of object complement constructions is to assign a proposition to the mental space of a subject (e.g., Givón 1993, Verhagen 2001, 2005). In relations like (14), a causal relation is embedded in a subject's mental space. In (17), on the other hand, a reason is given for a mental space plus its content. This difference is illustrated by Figures 1 and 2. Determining whether a relation holds between two clauses within a clausal complement or between one segment containing a clausal complement and another segment can be done by considering the mental representation of the discourse and determining between which units of thought the relation holds. This is an interpretation process, in which annotation and segmentation are mixed.

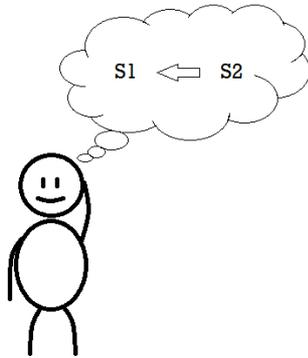


Figure 1. Coherence relation embedded in a mental space.

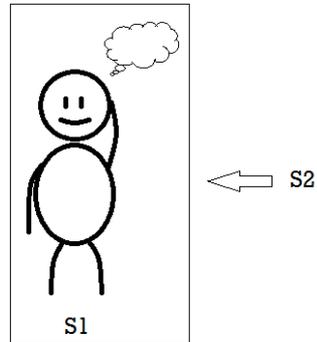


Figure 2. Coherence relation between a proposition and a proposition embedded in a mental space

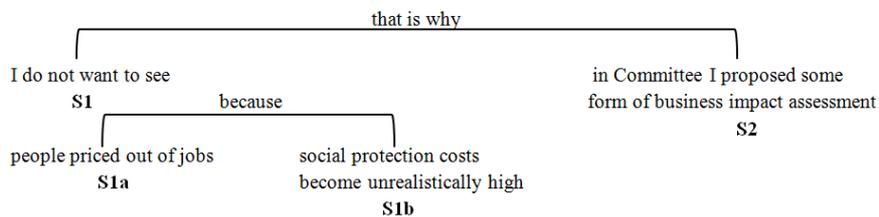
If we allow the segmentation of clauses within a complement, the entire relation can be treated as part of the main clause when considering the larger discourse structure. In (18), for instance, a causal relation holds between the two clauses of the complement of the verb *see*, as indicated in (18a). The next clause, *in committee I proposed some form of business impact assessment* appears to be a result, explicitly signaled by *that is why*, of the preceding main clause including its complements: because the speaker did not want people to lose their jobs over social protection costs, he proposed investigating the effects the social protection plans would have on businesses. This relation is segmented in (18b).

- (18) I am in favour of social protection, I am in favour of the original Commission document, but I do not want to see people priced out of jobs because social protection costs become unrealistically high. That is why in committee I proposed some form of business impact assessment, so that costs and risks to jobs could be taken into account, and the EPP-DE Group supported this amendment. {ep-00-02-15}
- (18a) ... I do not want to see [people priced out of jobs]_{S1} because [social protection costs become unrealistically high.]_{S2}
- (18b) ... [I do not want to see people priced out of jobs because social protection costs become unrealistically high.]_{S1} That is why [in committee I proposed some form of business impact assessment]_{S2} ...
- (18c) ... [I do not want to see [people priced out of jobs]_{S1a} because [social protection costs become unrealistically high.]_{S1b}]_{S1} That is why [in committee I proposed some form of business impact assessment]_{S2} ...

The fact that the main clause is stranded and not included in the first segment of the relation in (18a), as would be done when following Schilperoord and Verhagen's

(1998) method, has the advantage that it can be connected to other parts of the discourse without giving the impression that it features as the first segment of multiple relations. In addition, by segmenting the two relations as in (18a) and (18b), or as in the merged version in (18c), the segments of both relations accurately capture the idea units that are related to each other. Schematically, the segmentation in (18c) can be represented as in (19).

(19)



Because there is no relation on the discourse level between the main clause and the complement, there is no coherence relation indicated to hold between S₂ and S_{1a-1b} in (19). In (18c), there is no closing square bracket after *see*, indicating that even though a new segment featuring in another coherence relation begins, *I do not want to see* is by itself not an independent discourse segment.

In the approach to segmenting fragments containing complement clauses we proposed in this section, we have adopted Schilperoord and Verhagen's (1998) notion of conceptual dependency, but changed the way in which we apply this notion to discourse segmentation. While they proposed the possibility of including only the first clause of a complement in its host clause to conceptually complete the main clause and allowing additional complement clauses to be independent discourse segments, we argued in favor of also having the option of segmenting the clauses within a complement. The entire coherence relation can then be used to conceptually complete the main clause. This approach, unlike Schilperoord and Verhagen's (1998), allows us to distinguish fragments in which a clause is related to a preceding complement from fragments in which a clause is related to a preceding main clause containing a complement construction, i.e., (14) versus (17). In the next sections, we will demonstrate that the same segmentation approach can be applied to fragments with clausal subjects, restrictive adverbial clauses, or restrictive relative clauses.

2.3.2 Clausal subjects, restrictive adverbial clauses, and restrictive relative clauses

Clausal subjects and restrictive adverbial clauses are similar to clausal complements, since in all these constructions, the main clause is conceptually dependent on the

embedded clause. The segmentation approach proposed for text fragments with clausal complements can also be applied to fragments containing clausal subjects or restrictive adverbial clauses. (20) is an example of a sentence in which the subject is made up of two coordinated non-finite clauses. Both *releasing terrorist prisoners* and *seeking to buy them off with places in rigged government* meet the structural criteria of discourse segments, as discussed in Section 2.2: they contain predicates, and they are not coordinated within a phrase. These clauses can therefore be segmented as in (20). The additive coherence relation as a whole can function as the subject of the sentence.

- (20) [[Releasing terrorist prisoners]_{S1a} and [seeking to buy them off with places in rigged government]_{S1b} is exalting terrorism]_{S1} and [not eliminating it.]_{S2} {ep-01-09-05}

There will probably not be many fragments in which it is unclear whether only one clause or multiple clauses should be included in the clausal subject. The only scenario in which this could happen is if the sentence has a preceding coordinating sentence that could end with a non-finite clause.

For restrictive adverbial clauses there seems more room for ambiguity. Taking into account the propositional content of the segments in addition to the structural properties of the fragment, is therefore important. (21) and (22) have identical surface structures when considering the parts containing the restrictive adverbial clause. However, on the basis of the propositional content of the clauses, it can be determined that in (21) the final clause is part of the restrictive adverbial clause, while in (22) the final clause is connected to the preceding main clause, including its restrictive adverbial clause. The different segmentations of (21) and (22) reflect the differences in discourse structure between the two fragments. Note that Schilperoord and Verhagen (1998) would allow only the segmentation option in (22) (the fragment in (21) would be segmented in much the same way as (22), with *until* and the first clause following it included in the main clause, and the next clause as a separate discourse segment).

- (21) We are tired of the linkage of various directives in this package, with the excuse that we cannot look for a review of the European Works Council Directive until [it has been further bedded-in]_{S1} and [the European company statute is in place.]_{S2} {ep-01-02-13}
- (22) [This cannot begin until there is a cessation of terrorism]_{S1} and [only yesterday there was another suicide bomb attack in Jerusalem.]_{S2} {ep-01-09-04}

Restrictive relative clauses are slightly different from clausal complements, clausal subjects, and restrictive adverbial clauses, since they do not seem to conceptually complete another clause, but rather a noun phrase, or referent. In contrast with restrictive relative clauses, non-restrictive relative clauses do not seem integral to the conceptualization of referents, and have traditionally been regarded as discourse segments. The segmentation of the relation signaled by *if* in (23) will therefore be allowed in most annotation approaches.

- (23) *In addition, the Commission is now considering possible measures in the fields of technical assistance and trade, [which could be gradually extended]_{S1} if [North Korea makes progress in the areas I have mentioned.]_{S2} {ep-01-01-17}*

It seems, however, also possible for restrictive relative clauses to contain multiple clauses between which a coherence relation holds. In (24), in which dashes indicate clause boundaries, the noun phrase *the worried elderly people* is followed by a restrictive relative clause.

- (24) But on the BBC we saw a film recently // showing the deformed children and animals and the worried elderly people // who have decided to go back // because that [Chernobyl] was their home, // even though there is a risk. {ep 96-04-17}

The segmentation strategy Schilperoord and Verhagen (1998) propose for clausal complements does not appear to be equally applicable to restrictive relative clauses. While it is technically possible to consider only the first clause after *people* as the relative clause, this results in a conceptually incomplete referent, since the group of people denoted in this fragment is more detailed than *the worried elderly people who have decided to go back*. In addition, including only the first clause of the restricted relative clause in the first segment, as in (24a), results in a segmentation that does not accurately represent the units of thought related to each other. The fact that Chernobyl was home to many people is *not* the reason why the speaker saw a film on the BBC.

- (24a) # But [on the BBC we saw a film recently showing the deformed children and animals and the worried elderly people who have decided to go back]_{S1} because [that was their home]_{S2} ...

Again, a good way of arriving at a segmentation that represents the discourse structure without losing information is to allow segmentation within the embedded clause. This way, the segmentation does not only capture the causal relation within the relative

clause, as indicated in (24b), but also the coherence relation signaled by *even though*, as indicated in (24c).

- (24b) ... the worried elderly people [who have decided to go back]_{S1} because [that was their home,]_{S2} even though there is a risk.
- (24c) ... the worried elderly people [[who have decided to go back]_{S1a} because [that was their home,]_{S1b}]_{S1} even though [there is a risk.]_{S2}

Text fragments that contain embedded clauses are prone to have multiple possible interpretations, since clauses adjacent to an embedded clause can be related to either the embedded clause or another clause in the discourse, usually a main clause. As argued above, interpretations can be differentiated by means of segmentation if we allow embedded clauses to potentially receive the status of discourse segments. In order to arrive at a segmentation that accurately reflects the inferred discourse structure, it seems important and perhaps even unavoidable to take into account the propositional content of the clauses when segmenting texts.

2.4 Stance markers and discourse segmentation

All of the fragments presented in Section 2.3 contained embedded structures. We demonstrated that by segmenting embedded clauses and allowing them to conceptually complete their superordinate structures, it is possible to arrive at segmentation options that accurately represent the discourse structure and leave no elements unaccounted for. For complement constructions such as the ones following *because* in (25) and (26), however, this option is not available.

- (25) I would like to put it to the Commissioner that [she lost the battle with her colleague Sir Leon Brittan on this]_{S1} because [we understand that he is not very enthusiastic about dealing with the Norwegians and does not want to introduce restrictions.]_{S2} *He is frightened it might cause problems under the EEA agreement while we in the Committee on Fisheries and many people in Parliament take a different view.*
{ep-97-01-16}
- (26) I have now been informed that [the Council will not deal with my question or ten other Members' questions]_{S1} because [it claims it has not had time to prepare its replies.]_{S2} *I do not think that is acceptable.*
{ep-02-04-10}

In (25) and (26), *we understand that* and *it claims* appear to not be part of the idea units related by *because*. In (25) it is not the speaker's understanding of Sir Leon Brittan's dislike of Norwegians that caused the Commissioner to lose her battle, but

rather Brittan's dislike of Norwegians itself. Similarly, in (26) it is not the Council's claim it did not have time to prepare replies that leads to the speaker's questions not being dealt with, but rather the Council's (supposed) lack of time. In these fragments it appears that the first segments relate to only the complements of the clauses following the connective; the only function of the superordinate clauses *we understand that* and *it claims* seems to be to modify the content of the complement clauses (a similar fragment can be found in PDTB Research Group 2007:42, ex. 152). As was illustrated in Section 2.3, it is possible to leave initial matrix clauses outside the coherence relations, to have the entire coherence relations fall under their scope, and to connect the main clause, including its complement, to other parts of the discourse. Applying this approach to the second segments in (25) and (26), seems more problematic. First of all, the coherence relations would be indicated to hold between two units embedded under two different clauses. In addition, *we understand that* and *it claims* would be truly stranded. They cannot function as independent discourse units and are not related to other parts of the discourse. This would go against Mann and Thompson's (1988) criterion that all elements of a text should be included in the segmentation of that text. Yet, not excluding the superordinate clauses from the segments would go against our principle that discourse segments should represent the idea units that are related to each other, since *we understand that* and *it claims* do not seem to have a function within the coherence relations.

In this section, we will draw a parallel between fragments such as the ones in (25) and (26) and relations that contain stance adverbials and argue that discourse elements expressing stance can either have a function in the coherence relation as a whole, or merely modify one of the segments. After proposing that only the elements in a text that are part of the propositional content should obligatorily be included in the segmentation, we will present a solution to the segmentation problem fragments containing stance markers and complement-taking predicates represent.

2.4.1 Complement-taking predicates as stance markers

In Section 2.3 we focused mostly on the part of Schilperoord and Verhagen's (1998) conceptual dependency notion that stated that embedded clauses cannot enter into a relation with their host clause, but another aspect of the conceptual dependency criterion is that embedded clauses cannot enter into a relation with any other part of the discourse (that is not also embedded under the same structure). However, this does appear to be the case in the relations in (25) and (26), since only the complements of the predicates following *because* seem to make up the idea units related to the first segments. Potential explanations are that either the definition of conceptual dependency is faulty, or that the complement constructions following *because* in (25) and (26) are not typical instances of clause embedding. Here we will argue that indeed the latter may be the case.

Schilperoord and Verhagen's (1998) definition of conceptual dependency implies that subordinate clauses may be more important than their matrix clauses. When it comes to predicates with object complements in particular, there has been a lot of discussion about the exact nature of the relation between the complement and its host clause. Although analyses of complement-taking verbs differ slightly in their specifics, what they seem to have in common is that they consider the complement to be central to the proposition being expressed. Both Givón (1993) and Verhagen (2001, 2005), for instance, propose that object complement constructions assign some proposition, expressed in the complement, to (the mental space of) a subject, expressed in the host clause. Fetzer (2014:73) suggests that this aspect of complement-taking verbs makes them especially suitable to express epistemic stance about the proposition to which they are adjoined, since epistemic stance is "concerned with the speaker's evaluation of the certainty, possibility and probability of a state of affairs." Thompson (2002) even claims that complement-taking predicates (CTPs) are used to express epistemic stance, evidentiality, or evaluation in the majority of cases. Some complement-taking verb constructions, most of them with self-referencing subjects have grammaticalized and tend to be viewed as "parentheticals," the most notable example being *I think* (e.g., Aijmer 1997, Brinton 2008, Traugott 1995). These parentheticals are generally analyzed as epistemic stance markers modifying the content of the following clause (Fetzer 2014, Hunter 2016).

Given the observed parallel between epistemic stance markers and CTPs, it is worthwhile exploring whether in discourse segmentation CTPs can be treated the same as stance markers. This comparison seems especially justified given CTPs' ability to express not just epistemic stance, but other types of stance as well. Conrad and Biber (2000:57) identify three types of stance: *epistemic stance*, which comments on "the certainty (or doubt), reliability, or limitations of a proposition, including comments on the source of information," *attitudinal stance*, which conveys "the speaker's attitudes, feelings, or value judgements," and *style stance*, which describes "the manner in which the information is being presented."² It appears that CTPs can also express attitudinal and style stance. In (27), for instance, the CTP expresses attitudinal stance, since the speaker conveys his positive attitude toward the proposition in the embedded clause. In (28) the CTP comments on the form in which the embedded clause is presented, and is thus an example of style stance.³

² In this chapter we use Conrad and Biber's (2000) definition of epistemic stance, which includes evidentiality. Although we are aware of the ongoing debate on the exact relationship between evidentiality and epistemic stance (see e.g., Cornillie 2009 for an overview), we do not feel that this issue is crucial to the current discussion.

³ Note that the segmentation problem posed by text fragments containing CTPs cannot be solved by annotating *attribution*, be it as a coherence relation, as in SDRT or RST, or as another type of construction, as in PDTB. Neither the CTP in example (30) nor the one in example (31) fits the definition of an attribution relation, which is to indicate who is responsible for the information in a fragment. Still, these examples do exhibit the same scopal properties as CTPs that do encode attribution.

- (27) It is great that we are going to coordinate with the Americans.
{ep-00-06-14}
- (28) Let me just briefly reiterate that Parliament is provided in writing with a full list of the Commission's positions on each of the amendments.
{ep-02-10-22}

In the remainder of this section, we will demonstrate that stance adverbials can be part of the segments of a coherence relation, but can also occur outside of the relation, in which case they modify either the entire relation or one of the segments (Section 2.4.2). Subsequently, we will propose treating CTPs expressing stance in a way similar to adverbials of stance in discourse segmentation (Section 2.4.3).

2.4.2 Stance adverbials and segmentation

It seems possible to draw a parallel between (25) and (26), in which the second segments of the causal relations appear to be modified by their superordinate clauses, and relations in which S₂ is modified by a prototypical stance marker, for instance an adverbial, as in (29).

- (29) [I am glad that Commissioner Prodi is going to look at the EIB]_{S1} because, frankly, [that institution is inefficient and ineffective in aiding those firms which could be innovative and competitive if they just had that helping hand.]_{S2} {ep-03-03-26}

In (29) it is not the case that the speaker's being frank about the EIB's inefficiency is the reason for the speaker to be glad it is being investigated. *Frankly* does not play a role in the coherence relation, but seems to merely modify S₂. This is in contrast to relations such as the one in (30).

- (30) *They [transitory measures] are there for the time in which the market is still being directly regulated, but this whole package envisages a time when the entire market will operate under normal competition aspects.* [Those transitory measures should be clearly identified]_{S1} because [hopefully we will not need them in a few years' time.]_{S2}
{ep-01-06-12}

S₂ in (30) also has a clause-initial stance adverbial, but *hopefully*, unlike *frankly* in (29), does seem to be part of the coherence relation: the speaker's hope that transitory measures will not be necessary in the future is the reason for his stating that they should be identified.

Stance adverbials can also have scope over an entire coherence relation, in which case they resemble complement constructions such as the ones in (14) and (18).

Adverbials unequivocally have scope over a whole relation when they immediately precede the connective, as in (31a). Adverbials in other positions can also have scope over an entire relation: both (31b) and (31c) can, but need not, receive an interpretation similar to the relation in (31a).

- (31a) [The proportion of the complaints outside the mandate even increased slightly,] probably because [we received a growing number of complaints by e-mail.] {ep-00-07-06}
- (31b) Probably, [the proportion of the complaints outside the mandate increased slightly] because [we received a growing number of complaints by e-mail.]
- (31c) [The proportion of the complaints outside the mandate probably increased slightly] because [we received a growing number of complaints by e-mail.]

Determining whether adverbials are part of the idea units related to each other, as in (30), or whether their function is to modify one of the segments, as in (29), or the relation as a whole, as in (31), can be crucial for the annotation of the fragments. One of the features of coherence relations important in many annotation approaches is whether a relation holds in the real world (or a fictional world), or whether it is constructed in the speaker's mind. This distinction has received many labels over the years: content vs. epistemic and speech act (Sweetser 1990), semantic vs. pragmatic (Sanders et al. 1992), internal vs. external (Halliday and Hasan 1976), ideational vs. rhetorical (Mann & Thompson 1988, Redeker 1990), objective vs. subjective (Pander Maat & Sanders 2000, 2001), and others. Here, we will refer to this property of coherence relations as SOURCE OF COHERENCE, following Sanders et al. (1992). Certain adverbials can change a fact to a judgment, claim, or conclusion, e.g., *He is a judge*, vs. *He is probably a judge*, which can affect a relation's SOURCE OF COHERENCE and, consequently, the relation label ultimately attributed to a relation in annotation. Note that not all adverbials have potential consequences for annotation. Adverbials of time, for example, have the same scopal properties as other adverbials, but determining their scope will probably be less important in the process of discourse annotation than determining the scope of adverbials expressing stance.

2.4.3 Complement-taking verbs and discourse segmentation

Leaving an adverbial stranded, as in (29) and (31), seems less problematic than leaving an entire clause unaccounted for in the discourse structure, as in (25) and (26). However, CTPs and their complements do not always seem to correspond to typical host clause-embedded clause constructions, in which case the complement-taking predicate functions as a stance marker. Not incorporating a stance marker in the discourse structure seems acceptable, since stance markers are not part of the

propositional content of a text, but rather “the lexical and grammatical expression of attitudes, feelings, judgments, or commitment concerning the propositional content of a message” (Biber & Finegan 1989: 93).

If we adopt the view that CTPs can potentially function as stance markers, fragments like (25) and (26) immediately become less problematic. The relations are no longer supposed to hold between two clauses embedded under different structures, and the only elements not being part of the idea units are stance markers rather than content elements of propositions. Both *we understand that* and *it claims* are instances of epistemic stance: they mention the source of information, and, especially in (26), comment on the speaker’s idea of the actuality of the proposition.

The function of CTPs does not seem to be absolute. The same surface code, for instance *I know*, can have a different function depending on the context (Fetzer 2014). If CTPs can either express the mental space to which a proposition is assigned, or the speaker’s stance toward a proposition, it is crucial for the process of discourse segmentation to determine which one is the case. If the main function of a CTP is judged to be assigning a proposition to a mental space, the predicate should be accounted for in the discourse structure, since it is part of the propositional content of a text. If, however, a CTP is judged to function as a stance marker, it should be treated in a way similar to other stance markers, for instance adverbials. In that case, the CTP may be part of a segment, since a relation can be between a proposition including its stance and another segment, as we have shown in Section 2.4.2, but can also modify only one of the segments and be left out of discourse segmentation.

It should be noted that stance markers also function as mental space builders in that they open the speaker space (e.g., Dancygier & Sweetser 2012, J. Sanders & Redeker 1996). There is, however, a crucial difference in space building between CTP that function as stance markers and those that do not. If a CTP functions as a stance marker, the whole proposition, including the stance, is assigned to the mental space of the speaker. In (26), for instance, the status of *it has not had time to prepare its replies* is being questioned by the speaker. This process is different from the space building function of the CTP itself, which is to explicitly assign the contents of the complement to the mental space of the CTP’s subject, which may, but certainly need not be the speaker.

Determining the function of a CTP within a specific text fragment relies heavily on its context: the exact same surface structure can function as a stance marker in one instance, and only connect a proposition to a mental space in another. There are, however, a few characteristics that seem to increase or decrease the chances of a CTP being a stance marker. Cognitive verbs with a first person singular pronoun, such as *I think*, *I mean*, *I hope*, or *I believe* seem to function as stance markers more often than other cognitive verbs (e.g., Biber & Finegan 1989, Thompson 2002, Thompson & Mulac 1991). (32), for instance, is a colloquial example in which *I believe* functions as a stance marker: the speaker was not put in a small room because she believed there

were no other rooms left. Instead, a more accurate paraphrase seems to be that she received the small room because it was the only available room, or so she thinks.

(32) We got a small room because I believe it was the only one available.

Despite cognitive verbs with a first person singular subject being more likely to function as stance markers, cognitive verbs with a different subject can also mark stance, as (25) and (26) illustrate.

CTPs can occur with or without a complementizer. Some have proposed that having a zero complementizer is the grammaticalized form of CTPs, and that CTPs without a complementizer can function as stance markers, while the function of CTPs with a complementizer is to assign a proposition to a mental space (e.g., Aijmer 1997, Fetzer 2014). Others, however, propose that CTPs with complementizers can also function as stance markers (Kärkkäinen 2003, Thompson 2002). In addition, Kaltenböck (2009) argues that on the basis of prosody there is no reason to assume that a complementizer affects a CTP's status, i.e., main clause versus stance marker. The presence of a complementizer therefore does not seem to be a reliable basis for excluding the possibility of a CTP functioning as a stance marker, although it may increase the likelihood of the CTP assigning a proposition to a mental space (Thompson & Mulac 1991).

This section explained examples such as (25), (26), and (32), in which S_2 starts with a CTP that does not seem to function in the relation, by arguing that CTPs and their complements are not always host clause-subordinate clause constructions. Instead, the CTP can function as a stance marker, in which case it is not part of the propositional content of the segment, but rather modifies the propositional content of S_2 . Excluding a CTP from the representation of the discourse structure therefore seems justified when it functions as a stance marker, but when a CTP's main function is to assign a proposition to a mental space, it should be accounted for in discourse segmentation.

2.5 Discussion and conclusion

This chapter has presented a theoretical approach to text segmentation and argued that segmentation without interpretation does not always result in an accurate representation of the discourse structure. The issues addressed in this chapter were mainly illustrated by fragments taken from the Europarl corpus. This corpus consists of the written-out proceedings of the European Parliament, which consist of a combination of prepared and spontaneous speech and contains both monologue and dialogue. As such, Europarl is a highly hybrid corpus. Some of the problems addressed in this chapter may occur more often in written language, such as the complexity of some of the examples in Section 2.3, while other issues may be more

essential to speech. Stance marking, for instance, seems to be generally more frequent in spoken than in written discourse (e.g., Biber 2006, Conrad & Biber 2000), and the use of CTPs as stance markers in particular has also been claimed to be especially frequent in speech (Thompson 2002). So even though coherence relations with an S_2 modified by a CTP seem to be very rare in written discourse (to our knowledge these have not been discussed anywhere else, with the exception of one example mentioned in PDTB Research Group 2007), we expect them to be more often encountered in spoken discourse. Our proposal for dealing with CTPs in discourse segmentation, whether they are located in S_1 or S_2 , seems therefore particularly relevant now that discourse annotation is increasingly moving toward spoken and conversational data.

It should be noted that the account of complement-taking predicates in discourse presented in Section 2.4 focuses on English. While we believe that CTPs can function as stance markers in other languages as well, we question whether this fact alone always leads to constructions such as the ones in (25) and (26). When a CTP functions as a stance marker, the main clause has essentially become a function word, or discourse marker, while the subordinate clause functions as the main clause. This process appears to be mostly semantics-driven, since the basis seems to be the overlap in meaning between CTPs and other stance markers. English does not differentiate between main clauses and subordinate clauses in its word order or by any other means, which seems to enable such a change taking place. In languages that do syntactically distinguish main clauses from subordinate clauses, we do not expect to see discourse patterns similar to (25) and (26), since main clause/subordinate clause status is much more fixed. This, however, seems an issue worth exploring in future research.

Allowing embedded clauses to be segmented would lead to a more accurate representation of the structure of a discourse, but it would also increase transparency in discourse annotation, because the discourse segments will more accurately correspond to the units of thought that are inferred to be related to each other. If a fragment is, for example, segmented as in (17), partially repeated below, it can be assumed that the annotator interpreted the relation to hold between the main clause, including its embedded complement, and the clause following *because*. If, on the other hand, a fragment is segmented as in (14c) or (18a), both repeated below, it can be assumed that the annotator interpreted the relation to hold between the two clauses of the complement.

- (17) ... [it is an achievement that I am here tonight]_{S1} because [Air France cancelled my flight at 2.10 p.m.]_{S2} but I am here!
- (14c) [He may remember that [I complimented him]_{S1a} because [he had written an article in a journal complimenting Parliament on rescuing the internal market.]_{S1b}]_{S1}

- (18a) ... I do not want to see [people priced out of jobs]_{S1} because [social protection costs become unrealistically high.]_{S2}

The segmentation would unambiguously indicate between which units of thought annotators considered the relation to hold. In case annotators have not attributed the same relation label to a fragment, differences in segmentation would immediately pinpoint the source of disagreement between annotators.

Although incorporating interpretation in the segmentation process leads to more accurate text segmentation, it does pose a problem for automatic text segmentation, which is an important and promising technique being developed both within the discourse community and in the field of NLP research. By identifying specific contexts in which multiple segmentation options should be considered, we can limit the amount of text for which we have to take into account meaning during segmentation. While automatic text segmentation systems will not be able to disambiguate fragments, it would be possible for them to flag, for instance, complement constructions. Only the crucial parts of a text would then have to be manually checked by a post editor. As constructions with multiple segmentation options, this chapter pointed out complement constructions, restrictive relative clauses, restrictive adverbial clauses, or stance markers, but other linguistic contexts may also be identified as often being structurally ambiguous. Having an inventory of constructions that are especially prone to segmentation ambiguities can also help limit the amount of text for which meaning has to be taken into account in manual text segmentation. This would preserve the original concept of treating segmentation and annotation as two separate steps as much as possible.

This chapter has argued that while the grammatical clause is a functional basis for identifying discourse segments, it is sometimes necessary to take into account the propositional content of the text to arrive at a segmentation of a text that accurately represents the discourse structure and in which the discourse segments correspond to the units of thought related to each other. One of the segmentation issues where meaning can play a role is ellipsis, in which case the situation model can be taken into account to determine whether a structure is a clause with an ellipted subject and main verb, or rather coordinated nouns within a single phrase functioning as a direct object. We also argued in favor of amending Mann and Thompson's (1988) completeness constraint, i.e., the criterion that all elements should be included in the segmentation of a text, to pertain only to the propositional content of a discourse. Stance markers, which are not part of the propositional content of the text, may for instance be left out. Determining whether a stance marker should be included in a text segment, can be done by considering the interpretation of the text. Finally, we demonstrated that for fragments with embedded clauses, for instance clausal complements or relative clauses, multiple segmentation options should be considered. Using the interpretation of a text fragment can help to distinguish between distinct syntactic structures that

have identical surface structures, e.g., (17) versus (14c), and to arrive at an accurate representation of the discourse structure.

3 **Annotating discourse**

Using the Cognitive approach to Coherence Relations for discourse annotation

The Cognitive approach to Coherence Relations (Sanders, Spooren, & Noordman 1992) was originally proposed as a set of cognitively plausible primitives to order coherence relations, but is also increasingly used as a discourse annotation scheme. This chapter provides an overview of new CCR distinctions that have been proposed over the years, summarizes the most important discussions about the operationalization of the primitives, and introduces a new distinction (DISJUNCTION) to the taxonomy to improve the descriptive adequacy of CCR. In addition, it reflects on the use of the CCR as an annotation scheme in practice. The overall aim of the chapter is to provide an overview of state-of-the-art CCR for discourse annotation that can form, together with the original 1992 proposal, a comprehensive starting point for anyone interested in annotating discourse using CCR.

This chapter had been submitted as:

Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (*under review*). Using the Cognitive approach to Coherence Relations for discourse annotation.

3.1 Introduction

Annotating coherence relations refers to the process of attributing labels that best capture the relation inferred between two segments in a text to that relation. To annotate coherence relations, researchers make use of discourse annotation schemes. Discourse annotation schemes differ greatly in the number of relations they distinguish, ranging from two (Grosz & Sidner 1986) to 81 relations (Carlson & Marcu 2001). This is in part due to the fact that there is disagreement about how many distinct coherence relations language users actually infer and how specific these relations are. On the other hand, these differences seem to be caused by the varying purposes of the annotation schemes and the research traditions they originate from.

One approach to describing coherence relations that has been around for a while is the Cognitive approach to Coherence Relations (CCR; Sanders, Spooren, & Noordman 1992). Not originally designed as a discourse annotation approach, CCR defines four basic cognitive primitives that can be used to order the set of coherence relations language users infer between segments in a text. Since its introduction, CCR has primarily been used as a basis for experimental and acquisition research on discourse coherence; this research includes both studies aimed to verify the cognitive relevance of CCR's primitives and studies in which CCR's primitives are used as a point of departure for researching discourse coherence (see Sanders & Evers-Vermeul *in press* for an overview).

CCR is also increasingly used as a basis for discourse annotation. Using CCR as a discourse annotation can be appealing for several reasons. Since it consists of cognitively relevant primitives, CCR is applicable cross-linguistically. Indeed, it has successfully been used in discourse annotation projects covering several different languages: Dutch (e.g., Evers-Vermeul 2005, Spooren & Sanders 2008, Stukker 2005), English (see Chapter 4), German (Pit 2003), French (Pit 2003), Spanish (Santana, Spooren, Nieuwenhuijsen, & Sanders *submitted*), and Mandarin Chinese (Li, Evers-Vermeul, & Sanders 2016, Xiao, Li, Sanders, & Spooren *submitted*). In addition, CCR's primitives present a systematic approach to the categorization of coherence relations and have been shown to correspond to the distribution of connectives in various languages (e.g., Knott & Sanders 1998, Li 2014, Pit 2003, Sanders & Spooren 2013, Wei 2018). CCR's individual primitives also make it attainable to employ naive annotators in annotation projects; Scholman, Evers-Vermeul, and Sanders (2016) show that undergraduate students can use a step-wise version of CCR to produce decent quality annotations without extensive training. Not being entirely dependent on expert annotators helps cut down on time and expenses of traditional annotation projects and opens up the possibility of crowd-sourcing annotations. Furthermore, CCR's value combinations are often much more informative than end labels and can provide a better insight into annotator disagreements (Demberg, Asr, & Scholman 2017; see also Section 3.5.1). Finally, the

CCR taxonomy is easily applied to only a subset of relations. When, for instance, only considering coherence relations involving some form of contrast, it is clear which primitives and distinctions should be used during annotation; for approaches that use end labels this is not necessarily as clear.

There also appear to be some downsides to using CCR as a discourse annotation approach. Since CCR was designed to “identify the primitives in terms of which the set of coherence relations can be ordered,” it does not constitute a “complete descriptively adequate taxonomy of coherence relations” (Sanders et al. 1992:4). Since the original 1992 proposal, several additional distinctions have been proposed that aim to improve the descriptive adequacy of the taxonomy. However, these proposals are distributed over several individual papers. In addition, there appears to be some skewedness in how well the approach is developed for different types of relations; there has been a lot of debate on how to operationalize CCR’s primitives in the causal domain, but less so for other types of relations. In addition, several new distinctions have been proposed and frequently used within the domain of causal relations (e.g., VOLITIONALITY, PURPOSE), while fewer additional distinctions have been suggested for other types of relations.

This chapter provides an overview of new CCR distinctions that have been proposed over the years, summarizes the most important discussions about the operationalization of the primitives, and introduces a new distinction (DISJUNCTION) to the taxonomy to further improve the descriptive adequacy of CCR (Section 3.3). Finally, it reflects on the use of the CCR as an annotation scheme in practice in Section 3.4. The basic considerations of CCR and the original CCR taxonomy are first outlined in Section 3.2. Overall, this chapter thus provides an overview of state-of-the-art CCR for discourse annotation and forms, together with the original 1992 proposal, a comprehensive starting point for anyone interested in annotating discourse using CCR.

3.2 The Cognitive approach to Coherence Relations

The original CCR taxonomy was proposed in Sanders, Spooren, and Noordman (1992), and is very much in line with work by Hobbs (1978, 1979, 1990) and Kehler (1995, 2002), who also consider coherence relations to be cognitive entities and approach coherence relations by formulating a limited set of organizing principles. Sanders et al. (1992:2) define the concept of *coherence relation* as “an aspect of meaning of two or more discourse segments that cannot be described in terms of the meaning of the segments in isolation.” Coherence relations are the reason that “the meaning of two discourse segments is more than the sum of the parts” (Sanders et al. 1992:2). This basic property of coherence relations is referred to as the *relational surplus*; the criterion that CCR’s primitives have to be features of the relational surplus is the *relational criterion*.

In CCR, discourse relations are considered to hold between segments that are minimally clauses (e.g., Evers-Vermeul 2005, Sanders & van Wijk 1996); we will refer to this as the *clausal criterion*. The clausal criterion is closely related to the basic definition of coherence relations in CCR, since clauses are the smallest grammatical units that can function meaningfully in isolation (see Chapter 2 dissertation for a more elaborate discussion of the clausal criterion).

CCR considers coherence relations to be cognitive constructs. Its taxonomy is therefore intended to be cognitively plausible. For a distinction to meet the *cognitive plausibility criterion*, it should be observable in or make relevant predictions about language acquisition and language processing (Sanders et al. 1992). In addition, evidence for cognitive plausibility can be drawn from the system of linguistic markers. Knott and Dale (1994) argue that the distinctions made by connectives and cue phrases are indicative of the distinctions made in the minds of language users (see also Knott & Sanders 1998).

It should be noted that because CCR defines coherence relations as cognitive constructs, the labels attributed to coherence relations in annotation should correspond to the relation that holds in the mental representation of the discourse, i.e., the inferred relation. If a relation is marked by a connective or cue phrase in the text, it may well be the case that the annotated relation does not correspond to what is explicitly signaled by the linguistic marker. (1), for example, is marked by the connective *and*, but the relation that is inferred is a causal relation: the not marrying is interpreted as a consequence, albeit jokingly, of the chips-eating. (1) should thus be annotated as a causal relation, not as an additive relation as the connective might suggest.¹

- (1) *I would ask a man to open the bag for me — men open most containers for me — but then* [he would know I eat chips.]_{S1} *and* [he would never marry me.]_{S2}

In focusing primarily on the relations that hold in the mental representation of a discourse, CCR's approach to the depiction of coherence relations is distinctly different from 'bottom-up' annotation approaches that seem to place more focus on the linguistic markers of coherence relations, such as for example the Penn Discourse Treebank (PDTB: Prasad et al. 2008).

The original CCR primitives meet all three of CCR's criteria. They are properties of the relational surplus, thereby satisfying the relational criterion. They

¹ With the exception of a few simple relations we constructed ourselves for the sake of clarity, the vast majority of examples in this paper were extracted from actual utterances, from either fictional or non-fictional sources. We opted to use real examples to give a more realistic illustration of the type of coherence relations you would encounter in annotation tasks than simplified, prototypical examples would give. The examples were not collected systematically but rather selected because of their suitability to illustrate specific properties of coherence relations. The source for each example is provided in the appendix.

can also be used to describe relations that hold between clauses or larger discourse segments, thereby satisfying the clausal criterion. Finally, all primitives are cognitively plausible. The difference between *positive* and *negative* relations, which are distinguished from each other by the POLARITY primitive (see Section 3.2.1), can for instance be observed in processing (positive relations are processed faster than negative relations; Clark 1974, Murray 1997, Wason & Johnson-Laird 1971), language acquisition (positive relations are acquired earlier than negative relations; Bates 1976, Bloom, Lahey, Hood, Lifter, & Fiess 1980, Eisenberg 1980, Evers-Vermeul & Sanders 2009), and the linguistic system (positive and negative relations are prototypically signaled by different connectives). The remainder of this section will give an overview of the four original CCR primitives: POLARITY, BASIC OPERATION, SOURCE OF COHERENCE, and ORDER OF THE SEGMENTS.

3.2.1 POLARITY

Discourse relations hold between two propositions, expressed by S_1 , which refers to the first segment in the linear order of segments, and S_2 , which refers to the second segment. A relation with a *positive* value for POLARITY features P (antecedent) and Q (consequent), as in (2). A relation has a *negative* value for POLARITY if it features a negative counterpart of P, not-P, or Q, not-Q, as in (3).

- (2) [We liked Bob] $_{S_1}$ because [he was both different and apologetic.] $_{S_2}$
 (3) [They ... never failed to invite us to their houses] $_{S_1}$ although [they knew we would never come.] $_{S_2}$

In (2), S_1 presents a consequence (Q) of the cause (P) in S_2 . In (3), however, S_1 is a contrastive consequence (not-Q) of the cause (P) in S_2 ; a logical consequence of knowing someone never takes your offer could be to stop inviting them.

Positive relations are often expressed with connectives such as *and* or *because*. Negative relations are often signaled by connectives such as *but* or *although*. Although positive relations can often be turned into negative relations by negating one of the arguments, it should be noted that relations with a negative value for POLARITY do not necessarily contain lexical negation, as is illustrated by (4). Similarly, relations containing lexical negation can have a positive value for POLARITY, as can be seen in (5).

- (4) Although [it's inspired by the vinyl bars of Japan,] $_{S_1}$ [this spot chooses accessibility over authenticity.] $_{S_2}$
 (5) [I don't make them a lot] $_{S_1}$ because [I don't think it's fair to the other cookies.] $_{S_2}$

3.2.2 BASIC OPERATION

The category of BASIC OPERATION takes two values: *causal* and *additive*. A relation is *causal* if there is an implication relation between the two arguments ($P \rightarrow Q$), as in (6). Conditional relations, as in (7), also involve an implication relation and are categorized as having a causal BASIC OPERATION under the original CCR proposal.

- (6) [Phone service in the greater Chicago area was tied up for two hours Christmas Eve]_{S1} because [some kid called a phone-in show to get a wife for his father.]_{S2}
- (7) If [there was a fan club]_{S1} [I'd be the president.]_{S2}

A relation is *additive* if there is no causal relation between the segments and the only relation that can be inferred between the segments is $P \& Q$, as in (8).

- (8) [I'm worried]_{S1} and [I'm confused.]_{S2}

3.2.3 SOURCE OF COHERENCE

The main distinction made in the SOURCE OF COHERENCE of a discourse relation is between *objective* and *subjective*.² A discourse relation is objective when its two segments are related by their locutionary meaning; the relation is observable in the real world, as in (9).³ Subjective relations are related because of the illocutionary meaning of one or both of its segments; they involve the speaker's reasoning, as in (10); subjective discourse relations are often a reason or motivation for a claim or conclusion.

- (9) [A Harry Potter festival that was supposed to take place near Glasgow this summer has been cancelled,]_{S1} because [too many people wanted to go.]_{S2}
- (10) [Knitted gifts are great]_{S1} because [they are timeless and will last forever if taken proper care of.]_{S2}

A specific type of subjective relations are *speech act* relations. In a speech act relation one of the segments relates to the performance of the speech act in the other segment, for instance by offering a motivation or justification, as in (11), or by indicating the relevance of an utterance, as in (12). Speech act relations can also hold between two speech acts, as in (13).

² In the original 1992 proposal, the values of SOURCE OF COHERENCE were called *semantic* and *pragmatic*. These were later renamed as, respectively, *objective* and *subjective* in Pander Maat and Sanders (2000).

³ In using the term 'real world,' we do not only refer to the actual Earth, but also to the 'real world' in for instance fictional settings.

- (11) [How long are they going to take to cook?]_{S1} Because [you've got twelve minutes to go.]_{S2}
- (12) [There is a wonderful theatre program,]_{S1} if [she's interested in that.]_{S2}
- (13) [Why would it take an unusual woman to keep him company?]_{S1} And [why was he wearing a Russian astronaut on his lapel?]_{S2}

The SOURCE OF COHERENCE values are highly comparable to Sweetser's (1990) *domains of use*, with objective relations corresponding to Sweetser's *content* relations, and subjective relations including both *epistemic* and *speech act* relations. Other distinctions similar to CCR's SOURCE OF COHERENCE can be found in Halliday and Hasan (1976) and Martin (1992; internal vs. external), Redeker (1990; ideational vs. rhetorical), Mann and Thompson (1988; subject matter vs. presentational matter), Hovy and Maier (1995; ideational vs. interpersonal and textual), Pander Maat (2002; content vs. epistemic and interactional), and van Dijk (1977; semantic vs. pragmatic).

3.2.4 ORDER OF THE SEGMENTS

Discourse relations consist of (at least) two segments. The linearly first segment is always referred to as S₁; the linearly second segment is always S₂. The ORDER OF THE SEGMENTS feature refers to how P and Q of the BASIC OPERATION map onto S₁ and S₂. It takes two values: *basic* if S₁ expresses P and S₂ expresses Q, as in (14), and *non-basic* if S₁ expresses Q and S₂ expresses P, as in (15).

- (14) Because [they live in sub-tropical climates,]_{S1} [African penguins have to cope with both cooling down on land and keeping warm in the water.]_{S2}
- (15) [I had to talk loud]_{S1} because [the movie was loud!]_{S2}

The ORDER OF THE SEGMENTS is only relevant to causal relations, since additive relations are symmetrical in this respect.

3.3 Extensions of the original CCR taxonomy

POLARITY, BASIC OPERATION, SOURCE OF COHERENCE, and ORDER OF THE SEGMENTS are the original four primitives of the CCR taxonomy. Since the 1992 proposal, there has been a lot of discussion on how to operationalize the primitives, as well as proposals for new primitives or additional distinctions to CCR for discourse annotation. In this section, we provide an overview of the most important developments since the original CCR proposal.

It should be noted that many of the new distinctions (see Section 3.3) have been proposed for only a subset of relations. This does not necessarily mean that the same distinction could not also be annotated for other types of relations.

VOLITIONALITY (Section 3.3.1.2), for example, divides the subset of positive causal relations into *volitional* and *non-volitional* causal relations, a distinction that has been argued to be cognitively relevant on the basis of evidence from language processing, language acquisition, and linguistic systems. While VOLITIONALITY could also be annotated for, for instance, conditional relations, there is no clear evidence that the distinction between *volitional* and *non-volitional* is as cognitively relevant within the class of conditional relations as it is within the class of causal relations. If such evidence were to be found, the VOLITIONALITY distinction could easily be extended to apply to all implication relations with a positive value for POLARITY; the same holds for other distinctions as well. Limiting additional distinctions to apply to only those subsets for which there are indications that the distinction divides that class of relations into a cognitively plausible subset, helps to create a balance between descriptive adequacy on the one hand, and cognitive plausibility on the other.

3.3.1 Additional distinctions within original primitives

There have been proposals for additional distinctions within certain parts of the original CCR taxonomy. Unlike the original primitives, these additional distinctions apply only to a (small) subset of coherence relations. The proposed distinctions allow annotators to make more fine-grained contrasts, thus improving the descriptive adequacy of the CCR taxonomy. Additional distinctions have been proposed within the class of positive relations (TEMPORALITY), the class of positive objective causal relations (VOLITIONALITY and PURPOSE) and within the class of negative relations (DIRECTNESS).

3.3.1.1 TEMPORALITY

The CCR taxonomy has recently been proposed to be extended with a new distinction: TEMPORALITY. The original CCR proposal considers temporal relations to be a subtype of positive additive relations; Sanders et al. (1992:28) state that “the properties distinguishing temporal relations from other additive relations concern the referential meaning of the individual segments.” Temporality is thus taken to be a propositional, rather than a relational feature of coherence relations, and, as such, does not meet all the criteria necessary to be adopted into the CCR taxonomy. Evers-Vermeul, Hoek, and Scholman (2017), however, argue that TEMPORALITY does meet the relational criterion. They show that the temporal information in the propositional content of the segments is not always sufficient to establish a temporal coherence relation. In addition, they argue that the ordering of discourse segments in time can be only determined for a combination of the discourse segments; not for segments in isolation. As such, TEMPORALITY is a feature of the relational surplus and meets the relational criterion. TEMPORALITY is then argued to also meet all other CCR criteria. Temporal relations can hold between clauses, and the relevance of TEMPORALITY is

observable in language processing, language acquisition, and in the connective inventory of several different languages.

After discussing other options, Evers-Vermeul et al. (2017) argue that the best way of incorporating temporal relations in CCR is adding another primitive to the taxonomy. The proposed primitive distinguishes between relations that are ordered in time and relations that are not ordered in time. Positive additive relations that are ordered in time are relations that are most prototypically referred to as ‘temporal relations.’ As is shown in Figure 1, two additional steps make more fine-grained distinctions within the set of relations that are ordered in time: between sequential and synchronous relations and between sequential relations that are chronologically ordered and sequential relations that have an anti-chronological order. While not explicitly included in the original CCR taxonomy, the use of a ‘multi-level primitive’ that includes additional distinctions relevant to only a subset of relations is in line with later proposals for additional distinctions, such as VOLITIONALITY (see Section 3.3.1.2).

1	Temporal		Non-temporal
2	Sequential		Synchronous
3	Chronological	Anti-chronological	

Figure 1. The three-step temporality primitive

One benefit of having TEMPORALITY as a separate primitive is that two different types of order can be distinguished for causal and conditional relations: implication order, as depicted by the original ORDER OF THE SEGMENTS primitive, i.e., basic versus non-basic order, and temporal order, i.e. chronological versus anti-chronological order. These two orders will coincide for many relations, as in the positive subjective conditional relation in (16). The relation has a basic ORDER OF THE SEGMENTS, i.e., S₁ expresses P and S₂ expresses Q. It also has a chronological temporal order, i.e., the events expressed in S₁, saying something and it being verified as false, occurs before the event expressed in S₂, the conclusion that a lie was told. Sometimes, however, the two orders diverge, as in the positive subjective conditional relation in (17), which has basic order, but anti-chronological order, since the event expressed by S₁, avoiding the crowds, occurs after the event expressed by S₂, booking a trip during low-season. The idea that there is an underlying temporal order that is opposite from the ORDER OF THE SEGMENTS is underlined by the fact that the relation in (17) can be paraphrased as *you might want to book a trip here during an off-month, because **then** you can dodge the crowds*, while a similar construction cannot be used to paraphrase (16).

- (16) If [he says one thing and is verified the opposite,]_{S1} [that's a lie.]_{S2}

- (17) If [you want to dodge the crowds and wet T-shirt contests,]_{S1} [you might want to book a trip here during an off-month.]_{S2}

Although it was not explicitly addressed in Evers-Vermeul et al. (2017) whether the TEMPORALITY distinction is applicable to *all* coherence relations, we consider temporal order to be especially productive to relations with a positive value for POLARITY.

3.3.1.2 VOLITIONALITY

It has been proposed that within the class of positive objective causal relations, a distinction can be made between *volitional* and *non-volitional* relations (e.g., Pander Maat & Sanders 2000, Sanders et al. 1992, Stukker, Sanders, & Verhagen 2008; see also Mann & Thompson 1988). Volitional causal relations involve a thinking actor who is responsible for an event in the antecedent of the relation, as in (18), where the making event in S₁ is a volitional action. Non-volitional causal relations do not involve a volitional action. In the relation in (19), for example, the consequent does not involve an agent; one fact leads to the other. It should be noted that some languages have dedicated connectives for non-volitional causal relations, such as *daardoor* ‘that is why’ and *doordat* ‘because of the fact that’ in Dutch (e.g., Stukker et al. 2008).

- (18) [I make them a lot]_{S1} because [I have this indescribable need to constantly have new pillows.]_{S2}
- (19) [The game has changed]_{S1} because [the way we communicate has changed.]_{S2}

Pander Maat and Sanders (2000) propose that volitional causal relations have something in common with subjective causal relations (see Section 3.3.3). Both types of relations involve a Subject of Consciousness (SoC); a thinking entity involved in the relation. The main difference between volitional causal relations and subjective causal relations is that in subjective relations the SoC is involved in the *construal* of the relation (see Section 3.3.3.1), whereas in volitional causal relations, the SoC is not. Instead, the SoC in a volitional causal relation is usually an agent. In addition, the SoC in volitional relations is typically explicitly mentioned (onstage; see Section 3.3.3.2). While the speaker is responsible for the action in S₁ and the fact in S₂, the causal relation does not stem from the speaker’s mind and is observable in the real world. Non-volitional causal relations do not involve an SoC at all.

3.3.1.3 PURPOSE

Another distinction within the class of positive objective causal relations is PURPOSE (Sanders et al. *in press*). Purpose relations feature a volitional action for which the

motivation is an intended result. In (20), for instance, the adding of the smell is done to achieve the intended result of people knowing when there is a gas leak. Unlike the relation in (20), the relation in (21) does not feature an explicitly mentioned agent and instead uses a passive construction in S₁. While the relation in (20) has an explicit agent (*they*), the agent in (21) is implicit in the passive construction in S₁. Since the agent in (21) is not absent but merely unmentioned, it can still be classified as a positive, objective causal relation specified for PURPOSE. Several different languages have connectives that typically express PURPOSE relations, such as *so that* or *in order to* in English, or *zodat* ‘so that’ in Dutch.

- (20) *The gas is odorless, but [they add the smell]_{S1} so [you know when there's a leak.]_{S2}*
- (21) *[Services are being enhanced to remain open 24 hours]_{S1} so that [no one will have to stay on the streets during the cold snap.]_{S2}*

For causal relations specified for PURPOSE, determining the ORDER OF THE SEGMENTS is not entirely straightforward (e.g., Sanders et al. *in press*, Sanders et al. 1992). On the one hand, the relations in (20) and (21) are very similar to *result* relations (i.e., positive causal relations with basic order). On the other hand, they also bear similarities to volitional causal relations with a non-basic order like the one in (18), because the intended result is the motivation for executing the intentional action in the first place (see also Reese, Hunter, Asher, Denis, & Baldrige 2007:12-13). In CCR, the intended result in positive causal relations specified for PURPOSE should be considered the consequent, Q, while the volitional action should be considered the antecedent, P. The ORDER OF THE SEGMENTS in (20) and (21) is therefore basic.

3.3.1.4 DIRECTNESS

Pander Maat (1998) evaluates the original CCR taxonomy with respect to negative relations. He argues that the original primitive inventory is insufficient to capture all major distinctions between relations with a negative value for POLARITY. On the basis of a corpus annotation study and using linguistic evidence, primarily from the Dutch connective inventory, he proposes a new distinction to be applied to negative additive coherence relations: DIRECTNESS.⁴

Pander Maat (1998) poses that in negative additive relations, the two segments are compared to each other. This comparison is *direct* if “the propositions are

⁴ Pander Maat (1998) also discusses PERSPECTIVE (same perspective versus perspective change) as a potential new distinction for negative relations. However, this distinction does not appear to be completely productive across all negative relations and cannot be applied as systematically as all other CCR primitives and additional distinctions (see Pander Maat 1998:194, Figure 1). In addition, the PERSPECTIVE distinction mainly appears to be a property of the propositional content of the segments, rather than a relational feature. We will therefore not discuss the PERSPECTIVE distinction at length here.

themselves incompatible” (Pander Maat 1998:192); the propositional content of S_1 is in direct contrast to the propositional content of S_2 . The comparison can also be *indirect*, in which case the results or conclusions on the basis of propositions are incompatible. Direct, negative, objective, additive relations contain, for instance, a semantic contrast. In (22) the statements about Neilia and Jill are directly compared. In (23), on the other hand, an indirect, negative, objective, additive relation, it is not the segments themselves that are in contrast to each other, but rather the results of both segments (‘conflicting causal forces’); daily gains imply an improvement, but the second segment indicates a trend in the opposite direction.

- (22) [Neilia would always be Mommy,] S_1 but [Jill was Mom.] S_2
 (23) [Stock market notches daily gain,] S_1 but [posts largest weekly drop since early 2016] S_2

Within negative, subjective, additive relations, DIRECTNESS mainly distinguishes between *qualifications* and *concessions*. Sanders et al. (1992) categorize *concessions* as negative, subjective, additive relations. In their view (see also Spooren 1989), *concessions* are relations that feature two arguments in favor of opposing views, see Figure 2.⁵ *Concessions* are similar to relations with conflicting causal forces, as in (23), except for their SOURCE OF COHERENCE.

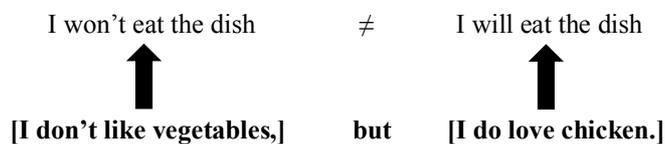


Figure 2. *Concession*

In *concessions*, the conclusion that can be drawn on the basis of the first segment is incompatible with the conclusion that can be drawn on the basis of the second segment. Since the causality is not found between the segments, but rather between the segments and their associated inferences, the relation between S_1 and S_2 is not an implication relation and, as such, the relation in Figure 1 is considered an additive relation. (24) and (25) are actual examples of *concessions*. In (24), the inference made on the basis of the first segment, “I won't agree with you,” is in contrast with the inference made on the basis of the second segment “I will agree with you.” In (25), the contrast holds between “you can write it yourself” and “we will have someone else write it.”

⁵ Outside of CCR, *concession* is also often used to refer to negative causal relations, e.g., “although she studied hard, she failed the exam.”

- (24) “This is a beautiful house.” “Thank you. I never know what to say when somebody says that. [You don’t want to agree]_{S1} but on the other hand, [it feels weird to disagree and say ‘no it’s a dump.’”]_{S2}
- (25) [I’m sure you would like to write the book yourself,]_{S1} but [your record is not what I might call promising, book-finishing-wise.]_{S2}

In *qualifications*, the second segment “cancels the strongest interpretation of the first statement” (Pander Maat 1998:186). As is illustrated in Figure 3, *qualifications* are similar to *concessions*, but the conclusion made on the basis of S₂ directly contrasts with the propositional content of S₁. While *concessions* are indirect, *qualifications* are thus direct.

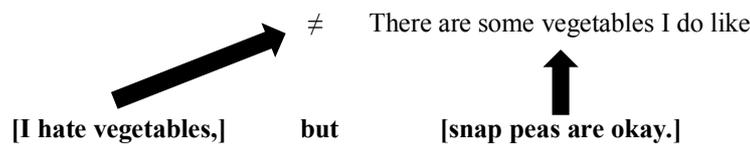


Figure 3. *Qualification*

(26) contains an actual example of a *qualification* relation. The proposition expressed in S₁, “I don’t know any blind people” is qualified by the statement that the speaker does know someone with a pretty severe eye condition, which implies that he does know someone who is *practically* blind.

- (26) [I, personally, don’t know any blind people,]_{S1} though [the guy I used to buy my newspaper from had pretty bad cataracts.]_{S2}

Pander Maat (1998) further distinguishes four specific types of *qualifications*: *simple qualification*, *exceptions*, *qualified denial*, and *denied intensification*. While taking note of these specific variations may be helpful in recognizing qualifications during annotations, the differences between the four types seem too fine-grained and segment-specific to be incorporated into the CCR taxonomy by means of additional distinctions (this is also not something Pander Maat (1998) proposes); they mainly refer to the direction of the *qualification* (weakening or intensifying) and to whether a stronger or weaker interpretation of the first segment should only be made to a certain extent or not at all.

Including the DIRECTNESS distinction within negative additive relations helps make the CCR taxonomy more descriptively accurate. In addition, Pander Maat (1998:199, Table 1) demonstrates that the differences between direct and indirect negative relations can be observed in the Dutch connective system, which suggests

that the distinction is also cognitively plausible. Finally, as Pander Maat (1998) points out, DIRECTNESS makes the CCR taxonomy more consistent, since the original 1992 proposal conflated SOURCE OF COHERENCE and DIRECTNESS for negative additive relations; the class of negative objective additive relations only included direct comparisons, while the class of negative subjective additive relations only included indirect comparisons.

3.3.2 Proposing a new distinction: DISJUNCTION

CCR was recently used as a tool to map other discourse annotation schemes onto each other (see Sanders et al. *in press*). The relation labels from RST, PDTB 2.0, and SDRT were ‘translated’ into CCR’s primitives, enabling a more accurate and straightforward comparison between the different frameworks than just comparing the end labels would have allowed. While CCR was able to capture the majority of distinctions, several extra features had to be formulated to ascribe a unique set of primitives and features to each relation label from a framework.⁶ Most extra features were similar to the distinctions discussed in Section 3.3.1 in that they were relevant to only a small subset of relations, and defined more specific instances of a certain relation type (e.g., LIST relations as a specific instance of positive additive relations). A notable exception was DISJUNCTION, a feature that distinguishes *disjunction* relations, in which the two segments are presented as alternatives, from other additive relations. Whereas RST, PDTB 2.0, and SDRT all include *disjunction* as a specific relation type, the original CCR taxonomy is unable to adequately capture the distinction between disjunctions and other types of additive relations.

3.3.2.1 Disjunctions in CCR

As the main reason for not including an “alternation relation” in their taxonomy, Sanders et al. (1992:29) refer to the “unclear status of *alternation*.” While some of the existing approaches to discourse coherence treated *disjunction* relations as a distinct class of relations, for instance “on a par with *conjoining*, *temporal*, and *implication*,” as Longacre (1983), others considered them a subcategory of *additive* relations (e.g., Halliday & Hasan 1976). In addition, as Sanders et al. (1992:29) point out, there was “also confusion about the nature of the *alternation* relation;” while some considered *disjunctions* to be primarily inclusive (e.g., Longacre 1983), others considered *disjunctions* to be primarily exclusive (e.g., Gamut 1982, Levinson 1983).

Here, we would like to argue in favor of including an additional distinction in the CCR taxonomy that can account for *disjunction* relations. Not only would such a distinction improve the descriptive adequacy of the taxonomy, it also seems to meet all criteria set by the CCR approach. First of all, *disjunction* relations hold between

⁶ Additional features were formulated if a distinction was made in at least two out of three frameworks. Note that these additional features were not proposed as new distinctions within CCR, but as necessary tools for the purposes of the Sanders et al. (*in press*) paper.

clauses, see (27), thereby satisfying the basic clausal criterion. DISJUNCTION is also a feature of the relational surplus, since the meaning of the relation as a whole is more specific than just the segments in isolation; without *disjunctions*, as in (27') the two segments would not be considered alternatives and both segments would be considered to be true.

- (27) [You either know it]_{S1} or [you don't]_{S2}
 (27') You know it // you don't know it

The final criterion that relational features have to meet before they can be included into the CCR taxonomy is cognitive plausibility. There seems to be ample linguistic evidence from connective inventories to suggest that DISJUNCTION is a cognitively plausible distinction, since many languages have connectives that prototypically mark *disjunctions*, for instance *or* or *either or* in English, *of* in Dutch, *oder* in German, *ou* in French, and *o* in Spanish. As discussed in Section 3.2, other evidence related to the cognitive plausibility of features of coherence relations can be derived from language acquisition and language processing. DISJUNCTION at the discourse level, however, does not seem to have received a lot of attention in these fields. A notable exception is a self-paced reading study by Staub and Clifton (2006). This experiment compares reading times of *disjunctions* in past tense and future tense signaled by *or* or *either or*. Staub and Clifton (2006) find that readers benefit more from the presence of *either* in the past tense condition than in the future tense condition. This suggests that when encountering a connective indicating DISJUNCTION after the first segment, readers have to update the truth-conditional status of S₁. This effect is much smaller, or even absent, in the future tense condition because the truth-conditional status of those segments is already uncertain. Staub and Clifton's (2006) experiment thus shows that DISJUNCTION can affect language processing and, as such, provides additional evidence in favor of the cognitive plausibility of DISJUNCTION.

3.3.2.2 DISJUNCTION as a new distinction in CCR

In line with the original Sanders et al. (1992) paper, we consider *disjunctions* to be a specific type of additive relations. Here, however, we propose to include DISJUNCTION as an additional distinction to the CCR taxonomy, applicable only to the class of additive relations. Similar to the additional distinctions discussion in Section 3.3.2, DISJUNCTION will carry the values *alternative*, in which case the segments are presented as alternatives, and *not alternative*, in which case the segments are not presented as alternatives. Additive relations that are *alternative* are the relations prototypically referred to as *disjunctions*; additive relations that are *not alternative* are all other types of additive relations.

As mentioned in Section 3.3.2.1, *disjunctions* can be exclusive, in which case the alternatives cannot hold at the same time, as in (27), or inclusive, in which case they can, as in (28).

- (28) [A little sweetener can take them from supper table to breakfast table]_{S1}
or [even turn them into dessert.]_{S2}

It is possible to distinguish between the inclusive and exclusive *disjunctions* using the POLARITY primitive (see also Sanders et al. *in press*). Since the two segments can hold at the same time, inclusive *disjunctions* have a positive value for POLARITY: P & Q. Exclusive *disjunctions*, on the other hand, always involve the negative counterpart of either P or Q: P & not-Q or not-P & Q. In (27), for instance, you know it, in which case you do *not* not know it, you do not know it, in which case you do *not* know it.

The SOURCE OF COHERENCE primitive applies to *disjunctions* as it does to other types of coherence relations; *disjunctions* can be either objective or subjective. In both (27) and (28), the alternatives are events that hold in the real world. As such, they have an objective value for SOURCE OF COHERENCE. In (29), both segments are opinions or claims, making the relation subjective. (30) is also subjective, since the *disjunction* holds between two speech acts, specifically between two questions.

- (29) Either [this person has lost her presence of mind]_{S1} or [she is just stupid.]_{S2}
(30) [Are you just feeling lazy]_{S1} or [do you need a break?]_{S2}

Since *disjunctions* are considered to be a subtype of additive relations, the ORDER OF THE SEGMENTS primitive does not apply. It should be noted that *disjunctions* are sometimes considered to include *unless*-relations (e.g., PDTB Research Group 2007, Reese et al. 2007; *unless you know it, you don't know it* has a meaning highly similar to the relation in [27]). In CCR, relations marked by *unless* are categorized as negative conditional relations; this also holds for relations not specifically marked by *unless* but with a similar interpretation.

3.3.3 Operationalizing SOURCE OF COHERENCE: segment-internal distinctions

The distinction between objective and subjective relations (or a similar distinction) is, as mentioned in Section 3.2.3, very common in theories about discourse and discourse annotation approaches. Although researchers seem to agree on prototypical examples, the SOURCE OF COHERENCE of a relation can be difficult to determine in the practice of actual corpus annotation (e.g., Sanders 1997). A proposal to improve the application of this primitive in the annotation of real-world examples is to make use of paraphrase tests, in which the segments of the relation are inserted in a paraphrase

that makes explicit either a subjective or objective reading, for instance *the fact that P causes S's claim/advice/ conclusion that Q* can be used to test whether positive causal relations with a basic order are subjective (Sanders 1997). Another practice that seems to facilitate determining the SOURCE OF COHERENCE of a relation is to consider the relation in its larger context, for example the whole text (Sanders 1997, Sanders & Spooren 2013).

It has been proposed that determining the SOURCE OF COHERENCE of a relation is difficult because while there are highly prototypical instances of objective and subjective relations, there are also many less prototypical examples (e.g., Degand & Sanders 1999, Sanders 1997, Stukker & Sanders 2012);⁷ non-prototypical examples are harder to classify than more prototypical examples. Several papers explore what makes a relation prototypically subjective or objective. Relevant features include the identity of the *subject of consciousness*, the explicit presence of the subject of consciousness, and the propositional attitude of the segments, each of which will be elaborated on in the rest of the section. Using these individual features can facilitate the process of determining a relation's SOURCE OF COHERENCE, as will be explained in Section 3.3.4.4. At the same time, the individual features are also used as additional distinctions within the SOURCE OF COHERENCE primitive to examine connective profiles in a more fine-grained way (e.g., Li 2014, Santana et al. *submitted*, Xiao et al. *submitted*).

3.3.3.1 Identity of the Subject of Consciousness

Pander Maat and Sanders (2000) propose that subjective relations involve a *Subject of Consciousness* (SoC) that is responsible for the construal of the relation; the relation stems from the SoC's mind (see also Pander Maat & Degand 2003, Pit 2003, Sanders, J. Sanders, & Sweetser 2009, J. Sanders, Sanders, & Sweetser 2012, among others). Subjective causal relations, for instance, involve the SoC's reasoning, as in (31). As was mentioned in Section 3.3.3.1, objective relations have either no SoC (non-volitional relations) or an SoC that is not responsible for the construal of the relation but is present as the agent of a volitional action (volitional relations).

- (31) [It must have been turkey mating season in Northern California]_{S1}
because [we've never seen so many turkeys strutting around.]_{S2}

In subjective coherence relations, the SoC is usually the speaker (Pander Maat & Sanders 2000): either the speaker or author of the discourse, as in (31), or the speaker responsible for the contents of a direct quote, as in (32). Alternatively, the SoC can be

⁷ Some have even claimed that it involves fitting a scalar phenomenon into distinct categories (e.g., Degand & Pander Maat 2003; Pander Maat & Degand 2001). See Stukker and Sanders (2012) for an overview of this argument, as well as Stukker and Sanders' argument in favor of a prototypicality account.

another actor in the discourse whose perspective is taken, as in (33). In (33), S₂ is a conclusion made on the basis of information in S₁. It does not say ‘so Tarzan concludes that the natives must be very near,’ but it is clear that the conclusion is drawn by Tarzan. Tarzan is the thinking entity responsible for the construal of the relation and therefore the SoC. In examples like (33), a third person actor temporarily becomes the speaker, although it would be even more accurate to say that in examples like these there is a ‘blend’ between the perspectives of the author or speaker and the discourse participant. (e.g., Sanders et al. 2009, J. Sanders & Spooren 1997).

- (32) “My intelligence can be very intimidating,” DeVos said. “And if [Donald Trump was a moron,]_{S1} [he would not want to be around people who are intelligenter than him]_{S2}.”
- (33) [Tarzan] was startled. Had he remained too long? Quickly he reached the doorway and peered down the village street toward the village gate. The natives were not yet in sight, though [he could plainly hear them approaching across the plantation.]_{S1} [They must be very near.]_{S2}

3.3.3.2 Explicit presence of the Subject of Consciousness

Not only the identity of the SoC, but also the extent to which the SoC is explicitly present in the relation has been argued to bear on the subjectivity of a relation. Langacker (1990, 1991, 2006) proposes that utterances with an explicitly mentioned, ‘onstage,’ speaker are more objective than utterances where the speaker is left implicit, or ‘offstage,’ since an explicitly mentioned speaker becomes itself the focus of attention. This view is applied to coherence relations by, for instance, Pit (2003), Sanders and Spooren (2015), and Stukker and Sanders (2012), who show that relations with onstage SoCs, as in (34), are less prototypically subjective than relations in which the SoC remains offstage, as in (34’). However, relations with an onstage speaker SoC do tend to be considered to be subjective relations if the relation is centered around a subjective judgment, opinion, or conclusion (e.g., Pander Maat & Degand 2001, Pander Maat & Sanders 2000, Pit 2003, Sanders & Evers-Vermeul *in press*, Wei 2018).

- (34) [I think all glitter should be banned,]_{S1} because [it’s microplastic.]_{S2}
 (34’) [All glitter should be banned,]_{S1} because [it’s microplastic.]_{S2}

It should be noted that a subjective relation can explicitly mention someone whose identity corresponds to the identity of the SoC, and still have an implicit SoC. In (35), for instance, the SoC is the speaker, but he is not explicitly mentioned in his role as SoC (as would be the case in *which I think was a bummer*). Instead, he is merely

explicitly mentioned as an actor in the event in S_2 that is used to motivate the judgment in S_1 .

- (35) *I made it through the night without getting fired.* [Which was a bummer]_{S1} because [I had spent the days previous applying for new serving jobs through Craigslist, just in case.]_{S2}

3.3.3.3 *Propositional attitude of the segments*

A final feature of coherence relations that is relevant to its SOURCE OF COHERENCE is the propositional attitude of the segments (e.g., Li 2014, Li et al. 2016, Sanders & Spooren 2009, 2015, Spooren & Degand 2010); are they, for instance, judgments, speech acts, or facts? Subjective relations prototypically involve judgments or speech acts, while objective relations prototypically feature facts. For implication relations, the propositional attitude of the consequent, Q, is most crucial (Li 2014).

3.3.3.4 *Determining SOURCE OF COHERENCE*

The identity of the SoC, the explicit presence of the SoC, and the propositional attitude of the segments can help in determining the SOURCE OF COHERENCE of a coherence relation. Perhaps the biggest advantage of using these individual features is that they can help uncover the source of potential annotation problems. If a relation is not prototypically subjective across the board, this may explain disagreements between annotators, or uncertainty within a single annotator. In addition, the individual features can be used to explicitly formulate a ‘cut-off point’ for categorizing relations as objective or subjective in an annotation project. This makes the SOURCE OF COHERENCE primitive easier to operationalize and makes the annotation process more transparent.

3.3.4 **State-of-the-art CCR for discourse annotation**

This section gave an overview of the most important developments in CCR since the original 1992 proposal when it comes to discourse annotation. Figure 4 provides a schematic overview of state-of-the-art CCR for discourse annotation. The overview is a flowchart resulting in unique value combinations at the bottom of the scheme. As is indicated by the grey shading and the prominence of POLARITY, BASIC OPERATION, and SOURCE OF COHERENCE, these are the only primitives relevant to all coherence relations. The distinctions in red squares are only relevant to the subset of relations below the primitive value to which they are attached. As such, they duplicate the set of relations below that primitive value. The numbers at the bottom of the scheme refer to the numbers in Table 1, where a simple, prototypical example is provided for each value combination in CCR. The segment-internal distinctions for SOURCE OF COHERENCE discussed in Section 3.3.3 are not explicitly incorporated in the scheme, but are considered to be part of the objective-subjective distinction within SOURCE OF

COHERENCE. For TEMPORALITY, we only included only the temporal order step for positive causal relations; by definition, these relations contain an underlying sequential temporal order.

Table 1.

Prototypical examples for each value combination in state-of-the-art CCR for discourse annotation.

1	positive causal objective basic chronological -volitional -purpose <i>Because it was raining, the streets were getting wet.</i>
2	positive causal objective basic chronological +volitional -purpose <i>Because it was raining, Jill brought her umbrella.</i>
3	positive causal objective basic chronological +volitional +purpose <i>Joe put up a tarp over the party area to prevent everyone from getting wet.</i>
4	positive causal objective non-basic anti-chronological -volitional -purpose <i>The streets were getting wet because it was raining.</i>
5	positive causal objective non-basic anti-chronological +volitional -purpose <i>Jill brought her umbrella because it was raining.</i>
6	positive causal objective non-basic anti-chronological +volitional +purpose <i>To prevent everyone from getting wet, Joe put up a tarp over the party area.</i>
7	positive causal subjective basic chronological <i>The streets are wet, so it must be raining.</i>
8	positive causal subjective basic anti-chronological <i>To prevent everyone from getting wet, you should cover the party area with a tarp.</i>
9	positive causal subjective non-basic chronological <i>You should cover the party area with a tarp to prevent everyone from getting wet.</i>
10	positive causal subjective non-basic anti-chronological <i>It must be raining, since the streets are wet.</i>
11	positive conditional objective basic chronological <i>If it rains, Jill will bring an umbrella.</i>
12	positive conditional objective non-basic anti-chronological <i>Jill will bring an umbrella if it rains.</i>
13	positive conditional subjective basic chronological <i>If Jill brought an umbrella, it must be raining.</i>
14	positive conditional subjective basic anti-chronological <i>If you want to prevent everyone from getting wet, you should cover the party area with a tarp.</i>
15	positive conditional subjective non-basic chronological <i>You should cover the party area with a tarp, if you want to prevent everyone from getting wet.</i>
16	positive conditional subjective non-basic anti-chronological <i>It must be raining, if Jill brought an umbrella.</i>
17	positive additive objective +temporal +sequence chronological -alternative <i>Joe put up a tarp before it started to rain.</i>

18	positive additive objective +temporal +sequence anti-chronological -alternative <i>Before it started to rain, Joe put up a tarp.</i>
19	positive additive objective +temporal +synchronous -alternative <i>While it was raining, Mona sat inside reading a book.</i>
20	positive additive objective -temporal -alternative <i>Mona read a book. She also built Legos with her son Mike.</i>
21	positive additive objective -temporal +alternative <i>Mike loves building with his Legos or even just taking his Lego creations apart.</i>
22	positive additive subjective -temporal -alternative <i>Legos are great. Reading is also wonderful.</i>
23	positive additive subjective -temporal +alternative <i>Jill is usually described as being great company or even as being someone who makes any party a success.</i>
24	negative causal objective basic <i>Even though it was raining, the streets stayed dry.</i>
25	negative causal objective non-basic <i>The streets stayed dry, even though it was raining.</i>
26	negative causal subjective basic <i>Even though it is raining, you should not bring an umbrella.</i>
27	negative causal subjective non-basic <i>You should not bring an umbrella, even though it is raining.</i>
28	negative conditional objective basic <i>Unless the skies have cleared, we are bringing an umbrella.</i>
29	negative conditional objective non-basic <i>We are bringing an umbrella, unless the skies have cleared.</i>
30	negative conditional subjective basic <i>Unless it is absolutely pouring down, you should not bring an umbrella.</i>
31	negative conditional subjective non-basic <i>You should not bring an umbrella, unless it is absolutely pouring down.</i>
32	negative additive objective direct -alternative <i>Jill brought an umbrella, but her friend did not.</i>
33	negative additive objective indirect -alternative <i>The rain is making the streets wet, but the sun is drying them really quickly.</i>
34	negative additive objective direct + alternative <i>The whole party it was either drizzling or pouring down.</i>
35	negative additive subjective direct +alternative <i>Every party last year was either really great or it was a total disaster.</i>
36	negative additive subjective direct -alternative <i>Rain is the absolute worst, though the smell of a light drizzle after a sunny day is pretty wonderful.</i>
37	negative additive subjective indirect -alternative <i>Going to that party sounds like fun, but it is pouring down outside.</i>

3.4 CCR as an annotation scheme in practice

In the introduction, we mentioned several advantages of using CCR for discourse annotation; it consists of cognitively plausible distinctions, is applicable cross-linguistically, and can be used by non-expert annotators. We also mentioned some potential problems researchers could run into when starting to use CCR, most of which we aim to solve in the current chapter. We provided an overview of all proposed additional primitives and distinctions and gave a summary of several discussions that have been carried out over separate research papers. This eliminates the need to sift through many different research papers to create an overview of state-of-the-art CCR. In addition, we took inventory of the full CCR taxonomy to see if there were any potential extra distinctions that would be eligible to be adopted into CCR and would increase the approach's descriptive adequacy. This led to our proposal for DISJUNCTION as a new distinction in CCR.

In this section, we reflect on the use of CCR as an annotation scheme in practice. To create the discourse-annotated parallel corpus that would serve as a basis for Chapters 4 and 5 of this dissertation, we annotated a large set of English coherence relations ($n \approx 2000$) extracted from the Europarl corpus (Koehn 2005). Here, we discuss several issues that are relevant to take into account when implementing the CCR taxonomy in a discourse annotation project: different options for calculating inter-annotator agreement, assumptions about the independence of primitives and distinctions and the possibility of also using end labels when using CCR, and some additional points concerning the operationalization of SOURCE OF COHERENCE.⁸

3.4.1 Calculating inter-annotator agreement when using CCR

When annotating coherence relations, researchers have to rely heavily on their own interpretation of the discourse, which is why discourse annotation is, at least to some extent, a subjective endeavor (e.g., Spooren & Degand 2010). To demonstrate that annotation has been done reliably and reproducibly, researchers can report an inter-annotator agreement measure: a (chance-corrected) numerical index that indicates the amount of agreement between two independent coders, such as Cohen's Kappa (Cohen 1960), Krippendorff's Alpha (Krippendorff 2004), Gwet's AC (Gwet 2002). For annotation efforts that make use of end labels to categorize coherence relations, the basis for calculating inter-annotator agreement is a confusion table like the one in Figure 5.

⁸ See Chapter 4 for a more extensive overview of the annotation project.

Coder 1	Coder 2			Total
	End label 1	End label 2	End label 3	
End label 1	agree	x	x	n
End label 2	x	agree	x	n
End label 3	x	x	agree	n
Total	n	n	n	N

Figure 5. Confusion matrix for annotation project using end labels

When annotating using CCR, calculating inter-annotator agreement is a little less straightforward. One option is to mimic the approach in Figure 5 and treat all primitive and distinction value combinations as end labels (e.g., ‘positive causal subjective non-basic,’ ‘negative additive subjective indirect’), but it is also possible to calculate agreement for each primitive or distinction individually (e.g., BASIC OPERATION: causal vs. additive). Treating all primitive and distinction value combinations as end labels has the main advantage that it makes the agreement scores comparable to other annotation efforts. However, the ‘end labels’ that are being compared are not entirely equivalent; relations have minimally three values (e.g., ‘positive additive objective’), but can have up to six values (e.g., ‘positive causal objective basic volitional purpose’). Calculating agreement for each primitive or distinction, on the other hand, is much easier than taking an ‘end label’ approach. In addition, it generates a clear overview of where exactly confusions or disagreements arise, which can be extremely valuable for further annotator training. Calculating agreement separately for each primitive or distinction makes it impossible, however, to check whether there is a systematic confusion between specific value combinations (e.g., ‘negative objective causal non-basic’ and ‘negative additive subjective indirect’), either because of annotator bias or because of a closer resemblance between two types of relations than the value combinations may suggest (see also Section 3.4.2). In addition, annotations can be dependent on the annotation of the other primitives or distinctions, especially when it comes to distinctions relevant to only a subset of relations. If one coder categorizes a relation as causal, while the other one marks it as additive, the two coders do not have the same number or type of other primitives and distinctions to annotate; coder 1 will for instance have to determine whether the relation is *CONDITIONAL*, while coder 2 has to make a decision on the *DISJUNCTION* distinction (see also Scholman et al. 2016 on the interdependence of annotations in CCR).

When using the full CCR taxonomy in an annotation project, it thus seems worth exploring the inter-annotator agreement both from the perspective of value combinations and for each individual primitive and distinction separately. The combination of both approaches will provide the most informative overview of annotations, as is also illustrated by the remainder of this section. When calculating inter-annotator agreement scores, it should be considered whether the annotation process, as well as the configuration in which they are being analyzed, match the

assumptions of the inter-annotator agreement statistic used; an inter-annotator agreement statistic may be unequipped to be used for annotations that are not independent or annotations that involve an uneven number of steps.⁹

3.4.2 Independence of primitives and the use of end labels in addition to primitives

While CCR's primitives are formulated as separate features, in practice the primitives seem to be slightly less independent than they may seem on the basis of the original taxonomy. First of all, the exact operationalization of a specific primitive or distinction can vary depending on other primitive values. As will be elaborated on in Section 3.4.3, determining the SOURCE OF COHERENCE for conditional relations involves a frequent problem that is much less often encountered in other types of relations: distinguishing between subjectivity and truth-value. In addition, agreeing on the BASIC OPERATION of relations with a positive value for POLARITY tends to be much easier than determining the BASIC OPERATION of negative relations; distinguishing between positive additive and positive causal relations is simple compared to distinguishing between negative additive and negative causal relations.

Another indication that primitives may not always be entirely independent from each other is that annotations may reveal a relatively frequent confusion between two types of relations that differ in multiple values. Based on the taxonomy, disagreement between relations that differ in only one value seems much more likely, and this type of confusion was indeed the most frequent type of disagreement in the annotation of the English relations in the parallel corpus (88% of all disagreements). The most common exception was a disagreement between annotators in which one annotator coded the relation as *negative causal objective, (non-)basic*, while the other coded it as *negative additive subjective indirect*, or vice versa.¹⁰ An example of such a relation can be found in (36). On the one hand, this relation could be analyzed as a negative objective causal relation, since setting targets and deadlines could plausibly lead to those targets and deadlines being met; the relation in (36) could then be analyzed as P leading to not-Q. On the other hand, the relation could also be analyzed as a negative subjective indirect additive relation (*concession*); the conclusion that can be drawn on the basis of S₁ is “we are doing great,” while the conclusion that can be drawn on the basis of S₂ is “we are not doing so great.”

⁹ Discussing the basic assumptions of commonly used inter-annotator agreement statistics and relating them to the possible ways in which CCR annotations could be analyzed is beyond the scope of this chapter, but see for instance Zhao, Liu, and Deng (2013) for a comprehensive overview of the basic assumptions of many inter-annotator agreement statistics. For a more general discussion on inter-annotator agreement in discourse annotation, see for instance Spooren and Degand (2010) or Hoek and Scholman (2017).

¹⁰ Distinguishing between negative additive and negative causal relations has also been reported as difficult or problematic on the basis of other annotation projects (e.g., Robaldo & Miltsakaki 2014, Zufferey & Degand 2017).

- (36) *We learned from that programme that implementation was not good enough. We have a solid base of more than 200 legal acts in the environment.* [We already have ambitious targets and deadlines in programmes,]_{S1} but [they have not all been met.]_{S2} {ep-01-05-30}

In practice, it can sometimes be harder to distinguish between two types of coherence relations than would be expected on the basis of the primitive and distinction value combinations in the CCR taxonomy. The observation that in practice, primitives are slightly less independent than they may seem to be in the taxonomy makes comparing annotations between coders using value combinations worthwhile. In addition, it makes it useful to explore the operationalization of a specific primitive or distinction within a specific subset of relations, e.g., TEMPORALITY within causal relations versus additive relations, or SOURCE OF COHERENCE within conditional versus causal versus additive relations.

Another possible solution is to use end labels in addition to the primitive value combinations. Some types of relations, especially highly specific types of relations, seem to become easier to recognize after becoming more familiar with relations that carry that specific combination of primitive and distinction values. It is for example very likely that inter-annotator agreement on relations with a negative value for POLARITY can be improved more by focusing on the exact difference between *qualifications* (negative subjective additive direct; see Section 3.3.1.4) and *concessions* (negative additive subjective indirect; see Section 3.3.1.4) than by further discussing the individual primitives.

Occasionally, it may thus seem easier to use end labels during annotation than individual primitives and distinctions. When encountering the relation in (37) in an annotation project using PDTB 2.0 (PDTB Research Group 2007), the relation label that should be chosen is fairly straightforward: *exception*. Using CCR, however, determining that (37) is a negative objective additive relation is, by comparison, much less obvious. Similarly, attributing a label to a relation like the one in (38) when using Carlson and Marcu's (2001) version of RST is simple: *otherwise*. Arriving at an annotation in CCR is much more involved: a negative objective conditional relation with basic order.

- (37) *Don't let the internet fool you — making hard boiled eggs in the microwave oven is trouble. If you try to hard boil eggs in your microwave you're likely to end up with a big mess to clean up. The rapid heat from the microwaves creates a lot of steam in the egg.*
[The steam has nowhere to go]_{S1} except [to explode out.]_{S2}

- (38) [When adding wine to a sauce, make sure you allow most of the alcohol to cook off;]_{S1} otherwise, [the sauce may have a harsh, slightly boozy taste.]_{S2}

However, differences in how easy it is to annotate certain types of relations exist not just between CCR and annotation approaches with end labels, but between annotation approaches in general. (37) is simple to categorize using PDTB 2.0, but is much harder to label using Carlson and Marcu's (2001) version of RST; (38) is straightforwardly labeled using Carlson and Marcu's (2001) annotation scheme, but much more difficult to categorize using PDTB 2.0. In sum, it can be worthwhile exploring which distinctions can be more reliably made when using end labels in addition to the individual primitives when using CCR for discourse annotation.

Another benefit of using end labels to refer to specific combinations of primitive values is that end labels can make talking about specific relation types much more convenient. It is for instance much easier to talk about *result* relations than to repeatedly mention 'positive objective basic order causal relations.' In such situations, the most obvious solution would be to define a relation type in terms of CCR primitives and distinctions and give it a single name to refer to the specific relation type. We took this approach ourselves in Section 3.3.1.4 of this chapter, where we used *qualification* to refer to negative additive subjective direct relations and *concession* to refer to negative additive subjective indirect relations. CCR's primitive approach is thus not incompatible with the use of end labels. The original CCR proposal by Sanders et al. (1992) already gives an overview of possible end labels that can be used to refer to specific combinations of primitive values. Being aware of which specific value combinations correspond to which type of end labels also makes it easier to compare CCR to other discourse annotation approaches and existing literature on coherence relations.

3.4.3 SOURCE OF COHERENCE versus truth-value

A common source of confusion pertains to the relationship between SOURCE OF COHERENCE and truth-value. Objective relations are defined to hold between two events in the real world, but this does not mean that the relation that is established between the two segments is necessarily *true*. In (39), for instance, the relation signaled by *because* is a positive volitional objective causal relation in which an SoC performs a volitional action for a specific reason. The relation as a whole, however, is a conclusion by the speaker, as is also indicated by *so*; the speaker makes a conclusion or claim about the unfolding of events in the real world. While the relation between S_1 and S_2 in (39) is an objective causal relation, the relation between that relation and the rest of the discourse is subjective. In practice, it can sometimes be difficult to distinguish between the SOURCE OF COHERENCE of the relation you are annotating and the SOURCE OF COHERENCE at a higher discourse level.

(39) So, [you're really just apologizing]_{S1} because [you need my advice.]_{S2}

(40) is a fragment extracted from the Europarl corpus. The relation at the end of the fragment is highly similar to (39), although it is slightly more complicated.

(40) *My group will also support the amendment, which other colleagues and I have signed in the name of the Socialist Group, for the deletion of paragraph 4. Why do we do that? Not because we necessarily disagree with the Scientific Committee on the issue of whether radiation can be safe for foodstuffs, but because it is not the whole story. The question is, why is this being done? Is it really for the benefit of the consumer? Is it something for which there is a consumer demand? If that was the case, we would not have as many cases as there are, certainly in my country, of illegal and covert irradiation.*
 [This has been carried out]_{S1} because [they do not want consumers to know about it.]_{S2} {ep-02-12-16}

The final sentence of (40) is a positive volitional objective relation, since it holds between an intentional act and a reason for that act. However, from the fragment it is clear that the discourse relation is the speaker's answer to the question *why is this being done?* The relation is *claimed* to be true: the reason for the intentional act is invented or hypothesized by the speaker. This does not, however, mean that the relation itself becomes subjective. Internally, the way in which S₁ relates to S₂ is objective, and without context, there would probably be no confusion. The subjective nature of the final sentence in (40) arrives from, and can be captured by, it as a whole being a claim and part of a *subjective* relation; the speaker claims that it is being done not for the benefit for the consumer, not because consumers demand it, but rather because consumers are preferred to not know about it.

The distinction between SOURCE OF COHERENCE and truth-value seems especially relevant to conditional relations, since they often seem to entail speaker involvement. Conditionals, of which the content usually has not been realized, are often predictions. In the relation in (41), for example, the speaker announces what his party will do in a certain scenario. Similar to the relations in (39) and (40), the relation between the two segments in (41) is objective, while the relation as a whole is a prediction. Here too, the SOURCE OF COHERENCE within the relation is not the same as the SOURCE OF COHERENCE of the relations that holds at a higher discourse level, i.e., between the relation as a whole and the preceding discourse.

- (41) If [we find that any Member of this House or their employees collaborated with the BBC in this farrago]_{S1} [we will expose them to the opprobrium of this House.]_{S2} {ep-00-02-14}

Removing the conditionality from the discourse relation helps when annotating the SOURCE OF COHERENCE of a conditional relation; if the resulting causal relation is objective, the conditional relation is also objective. Without the conditionality, (41) would become ‘there has been a collaboration with the BBC, which is why we will expose them;’ a positive volitional objective causal relation. It is also not uncommon for conditional relations to express the speaker’s *negative stance* toward the antecedent, and therefore toward the entire prediction, actually taking place (sometimes also called *counterfactual* or *irrealis*). In English, indicating that something is unlikely to come true can for instance be done by means of a distanced verb form, e.g., *if we found that*. Speakers can also encode that the event did definitely *not* take place, e.g., *if we had found that*. Even though negative stance seems to emphasize the presence of a speaker, it does not usually influence the SOURCE OF COHERENCE between the segments of the conditional relation. With negative stance added, the relations in (41) for example still expresses that if one real world event occurs, it leads to another real-world event.

In general, an increased awareness of the difference between SOURCE OF COHERENCE and truth-value and about the way in which the SOURCE OF COHERENCE at a higher discourse level can influence the way in which the SOURCE OF COHERENCE between two segments is perceived can help improve the quality and reliability of discourse annotation.

3.5 Conclusion

The Cognitive approach to Coherence Relations was originally proposed as a set of cognitively plausible primitives to order coherence relations, but is also increasingly used as an annotation scheme for classifying coherence relations. In this chapter, we gave an overview the most important developments within CCR from the point of view of discourse annotation. We discussed proposals for new primitives and additional distinctions, and summarized the discussion on how to operationalize an original primitive, SOURCE OF COHERENCE. In addition, we argued in favor of adding a new distinction to CCR: DISJUNCTION. Finally, we discussed some practical issues we encountered during a recent annotation project using CCR. As a whole, this chapter gives an overview of state-of-the-art CCR for discourse annotation. As such, it can be used, together with the original 1992 proposal, as a point of departure for anyone interested in annotating coherence relations using the Cognitive approach to Coherence Relations.

4 Cognitive complexity and the linguistic marking of coherence relations

Coherence relations can be made linguistically explicit by means of connectives (e.g., but, because) or cue phrases (e.g., on the other hand, which is why), but can also be left implicit and conveyed through the juxtaposition of two clauses or sentences. However, it seems that not all relations are equally easy to reconstruct when they are implicit. In this chapter, we explore which features of coherence relations make them more, or less, likely to be conveyed implicitly. We adopt the assumption that expected relations are more often implicit than relations that are not expected, and propose to determine a relation's expectedness using the notion of cognitive complexity. We test our hypotheses by means of a parallel corpus study, in which we analyze the translations of explicit English coherence relations from the Europarl Direct corpus into four target languages: Dutch, German, French, and Spanish. We find that cognitive complexity indeed influences the linguistic marking of coherence relations, and that this does not vary between the languages in our corpus. In addition, we find that a relation's relational and syntactic dependency also influences its linguistic marking, but that these measures are not completely independent of relation type.

This chapter has been published as:

Hoek, Jet, Zufferey, Sandrine, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2017). Cognitive complexity and the linguistic marking of coherence relations: A parallel corpus study. *Journal of Pragmatics* 121, 113-131.

4.1 Introduction

Coherence relations such as cause-consequence, as in (1), and temporal sequence, as in (2), connect idea units in a discourse and can be defined as “an aspect of meaning of two or more discourse segments that cannot be described in terms of the meaning of the segments in isolation” (Sanders, Spooren, & Noordman 1992:2). Such discourse segments consist of minimally a clause.¹

- (1) Because [Peter had stolen some cash from his grandma’s purse,]_{S1} [he was grounded for a month.]_{S2}
- (2) [Jane and Harry did a few jumping jacks]_{S1} before [they started their morning run.]_{S2}

Coherence relations can be made linguistically explicit by means of connectives (e.g., *but*, *because*) or cue phrases (e.g., *on the other hand*, *which is why*), but can also be left implicit and conveyed through the simple juxtaposition of two clauses or sentences. In the absence of a connective, readers or listeners have to infer the relation between the discourse segments themselves using the linguistic context and their world knowledge (Kintsch 1998, Zwaan & Radvansky 1998). However, it seems that not all relations are equally easy to reconstruct when they are implicit. In (3), a fragment taken from the Europarl corpus (Koehn 2005), for instance, the relation between the first and the second sentence has not been explicitly marked by a connective, but it can still be determined that the potential hate-inducing qualities of the remarks are the reason for the speaker to find them unworthy of being uttered in the European Parliament; an appropriate connective would have been *because*. The two clauses that make up the second sentence, on the other hand, are connected by *if*. Leaving out this connective would make the relation hard to reconstruct; although the remarks would still be understood to cause agitation with the public, them getting into the media would most likely be interpreted as a given, rather than as a hypothetical event.

- (3) Those sort of remarks are unworthy of this Parliament. If [they get into the media]_{S1} [it just stirs up hate.]_{S2} {ep-97-10-23}

The intuition that some types of relations can be more easily left implicit than others is confirmed by analyses of discourse-annotated corpora. Studies by for instance Asr and Demberg (2012) on the PDTB (Penn Discourse Treebank; Prasad et al. 2008) and Taboada (2006) on a corpus annotated using RST (Rhetorical Structure Theory; Mann

¹ See Chapter 2 for a more detailed discussion of what defines a discourse segment. A discourse relation’s segments will be indicated by square brackets throughout the paper; the connective marking the relation is left outside the segments.

& Thompson 1988) show comparable patterns in the marking of coherence relations. Causal relations and relations of general addition, for instance, are often expressed implicitly. Conditional relations and relations involving concession, on the other hand, tend to be explicitly marked. While the existence of asymmetries in the marking of coherence relations has been clearly established, the exact mechanisms that cause these asymmetries are not yet fully understood.

The question of when a coherence relation has to be explicitly marked in order to avoid a loss of coherence is an important issue within the study of discourse, but it is relevant to other fields as well. For instance, whether or not coherence relations are explicitly marked has been found to influence the processing and comprehension of educational texts (e.g., McNamara, Kintsch, Donger, & Kintsch 1996, van Silfhout, Evers-Vermeul, & Sanders 2015). In addition, learning to appropriately mark coherence relations is a vital but difficult aspect of acquiring a second language, with L2 speakers regularly over- and underusing connectives, using them in non-prototypical constructions and contexts, and not always understanding when relations should be marked explicitly (e.g., Granger & Tyson 1996, Müller 2005, Zufferey & Gygax 2017). Furthermore, knowledge about the explicitness vs. implicitness of coherence relations can serve as valuable input for various NLP applications, such as automatic language generation and machine translation (Meyer & Popescu-Belis 2012). Finally, understanding what drives the marking of coherence relations can inform more general topics of how we process language and establish coherence within a discourse.

4.1.1 Expectedness and the marking of coherence relations

Asr and Demberg (2012) propose that the linguistic marking of a coherence relation is strongly influenced by a relation's *expectedness*, with expected relations being more often left implicit. This assumption finds its roots in the Uniform Information Density (UID) hypothesis (Frank & Jaeger 2008, Levy & Jaeger 2007), which proposes that speakers "structure their utterances so as to avoid peaks or troughs in information density" (Levy & Jaeger 2007:1). A linguistic element that marks something that was already expected by the reader hardly adds any information to the discourse, and therefore constitutes a trough in information density. Conversely, leaving implicit something that was not already projected causes an overload of information to be extracted from the linguistic elements that are present, thus constituting a peak in information density.

The idea that expected relations can be left unmarked can also be thought of in terms of effort versus effect, key notions from Relevance Theory (Sperber & Wilson 1985, Wilson & Sperber 2005). If an unexpected relation is not marked, its inference requires too much effort for the resulting cognitive effect. As a result, an easier, more expected coherence relation will be inferred. Explicitly marking unexpected relations therefore ensures that the right relation is established. For example, not explicitly

signaling a conditional relation, such as the one in (3), makes it hard or even impossible to recover that relation, and the fragment will most likely receive a different, non-hypothetical interpretation.

If we assume that the linguistic marking of coherence relations is to a large extent governed by expectedness, it is key to determine which kinds of relations are expected to occur in a discourse. Asr and Demberg (2012) refer to Murray's (1997:228) *continuity hypothesis* (see also Segal, Duchan, & Scott 1991), which states that "readers have a bias toward interpreting sentences in a narrative in a continuous [linear] manner," with for instance, additive and (forward) causal relations expressing continuity, and examples of discontinuity being "reversions to an earlier setting or scene (such as a flashback), an abrupt topic change, a surprising turn of events, a character moving away from what he/she is doing, or a violation of an expectation created in the previous text." In addition, Asr and Demberg (2012) take into account Sanders' (2005:9) *causality-by-default hypothesis*, which proposes that "because readers aim at building the most informative representation [of a text], they start out assuming the relation between two consecutive sentences is a causal relation." These two hypotheses predict that readers expect temporally forward relations, additive relations, causal relations, and relations that do not involve a form of contrast. However, they are not sufficient to account for all the observed differences in the marking of coherence relations in corpus data. For example, neither the continuity hypothesis nor the causality-by-default hypothesis is able to account for the fact that conditional relations tend to be explicitly marked, as these are neither non-causal nor reversing the temporal order of events (see also Asr & Demberg 2012:6). In addition, the predictions about the marking of coherence relations that can be formulated on the basis of the continuity hypothesis and the causality-by-default hypothesis are very coarse-grained and unable to make lower-level distinctions. For instance, it follows from both hypotheses that contrast relations, e.g., (4), as well as concessive relations, e.g., (5), are unexpected. However, these two relations seem to differ in the frequency with which they are explicitly marked, with contrast relations occurring implicitly more often than concessive relations (Asr & Demberg 2012:9).

- (4) [Frank made a respectable \$175 at the flea market,]_{S1} but [Kate went home with \$630.]_{S2}
- (5) Although [Jamie had never worked a day in his life,]_{S1} [he was insanely rich.]_{S2}

In this chapter, we further explore which features of coherence relations make them more, or less, likely to be conveyed implicitly. We adopt the assumption that expected relations are more often implicit than relations that are not expected, and propose to determine a relation's expectedness using the notion of cognitive complexity. Traxler, Bybee, & Pickering (1997) find that simple relations are

processed faster than complex relations. They argue that readers construct the simplest possible coherence relation and adapt their representation of the discourse if this relation is not consistent with the context (see also Traxler, Sanford, Aked, & Moxey 1997). This observation seems to be largely in line with the predictions made by the continuity hypothesis and the causality-by-default hypothesis, as will be discussed in Sections 4.2.4 and 4.2.3 respectively. If cognitively simple relations are expected, they should occur implicitly more often than cognitively more complex relations.

In discourse processing, the notion of expectation has been proven to be relevant in at least two ways (see for instance Renkema 2004 or Traxler & Gernsbacher 2011 for overviews). First, it has often been shown that readers have expectations about the upcoming text in terms of content: readers use their world knowledge, organized in frames, scripts or scenarios to predict the content of the rest of the text. For instance, in a fragment in which two people enter a place where somebody shows them where to sit, readers will infer this person is the waiter and predict that he will bring the menu soon. This knowledge belongs to a stereotypical restaurant scheme, and the schematic expectation steers discourse processing in a top-down way. A second way in which expectations can play such a role is in terms of structure: given a first part of text, readers have expectations about the structural role the next part will play. Such expectations differ with genres: in stories, one event will be followed by another, ordered in a temporal sequence; in a newspaper article a problem may be followed by a solution and in experimental articles a methodology section will follow the experimental hypotheses.

In this chapter, we operationalize expectedness in a way that is related to the second way of operationalizing expectation mentioned above; our operationalization also pertains to discourse structure, or, more specifically, the coherence relations that hold between idea units. However, it is also slightly different, since we focus on the relative cognitive complexity of relations. We assume that expectations based on the complexity of coherence relations constitute default, baseline expectations, on top of which more specific topical or structural predictions such as the ones mentioned above can be formed.

We adopt the Cognitive approach to Coherence Relations (CCR; Sanders et al. 1992, and later work) as our discourse annotation framework and as a way to determine cognitive complexity. CCR distinguishes four basic primitives of coherence relations. Determining the complexity of the primitives' values using evidence from processing and corpus-based research allows us to make fine-grained predictions about the linguistic marking of coherence relations, see Section 4.2.

4.1.2 Relational and syntactic dependency and the marking of coherence relations

In addition to relation type, other factors have been proposed to affect influencing the linguistic marking of coherence relations as well. Patterson and Kehler (2013), for

instance, take into account the position of the coherence relation in the larger discourse structure. The model they build to predict the marking of coherence relations uses information about whether the relation is embedded in another relation (Rel2 in Figure 1), whether it contains another relation (Rel1 in Figure 1), or whether it shares one of its segments with another relation (Figure 2). This *dependency* information was taken from the PDTB (Prasad et al. 2008), on which the model was trained and tested.

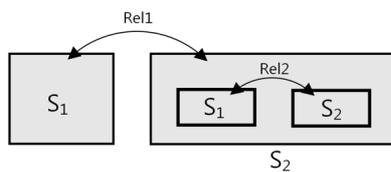


Figure 1. Embedded (Rel2) and containing (Rel1) relation

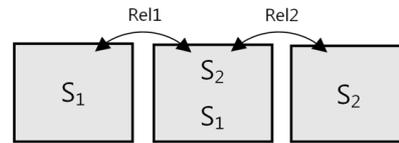


Figure 2. Shared segment

Patterson and Kehler (2013) find that embedded relations are often explicitly marked, and that containing another relation and sharing a segment with another relation both increase a relation's likelihood to be implicit. The current study also takes into account the dependency between coherence relations in the discourse (*relational dependency*).

In addition to relational dependency, we take into account each relation's *syntactic dependency*, for which we determine whether the relation is embedded under a syntactic construction, for instance a complement construction, as in (6), or a relative clause, as in (7).¹

- (6) *I have every confidence that if [we and the politicians of Europe support him]_{S1} [we will see the kind of reform we need.]_{S2} {ep-00-01-18}*
- (7) *Finally, I wish to speak particularly for my constituents in East Anglia, [many of whom have not drawn on the common agricultural policy]_{S1} because [they grow crops which have never been subsidised by this Union.]_{S2} {ep-02-06-11}*

We hypothesize that relations that are syntactically dependent will have a higher likelihood of being explicit than relations that are syntactically independent, on a par with the hypothesis for relations that are embedded in another coherence relation.

¹ For a discussion of coherence relations embedded under syntactic constructions and the discourse segment status of syntactically embedded clauses, see Chapter 2.

4.1.3 Outline

We test our hypotheses by means of a parallel corpus study, in which we analyze the translations of explicit English coherence relations from the Europarl corpus into four target languages: Dutch (NL), German (DE), French (FR), and Spanish (ES). Section 4.2 provides an overview of CCR, discusses the cognitive complexity of the primitives' values, and formulates hypotheses concerning the marking of coherence relations. Section 4.3 outlines the parallel corpus study. Results from this study are reported in Section 4.4 and discussed in Section 4.5.

4.2 Cognitive complexity of coherence relations

In this study, we use CCR as our discourse annotation framework and annotate coherence relations using four primitives: POLARITY, BASIC OPERATION, SOURCE OF COHERENCE, and ORDER OF THE SEGMENTS. In order to establish the complexity of each primitive's values, in this section we consider evidence from different areas of linguistic research. We consider logic as a first indicator of complexity. In addition, we look at results from processing studies. Slow processing implies a high processing cost, which is associated with cognitive complexity, i.e., simple relations are processed faster than difficult relations (Köhne & Demberg 2013 Traxler, Bybee, & Pickering 1997 Sanders et al. 1992, 1993). We also consider patterns in language acquisition. The order in which different types of coherence relations are acquired has been related to the complexity of the relations, with more complex relations being acquired later than simpler relations (e.g., Evers-Vermeul & Sanders 2009, 2011, Spooren & Sanders 2008, van Veen 2011). Finally, we use Mental Space theory (Fauconnier 1985) to assess a relation's complexity. Specifically, we take into account whether a relation involves a mental space shift. Mental spaces are roughly equivalent to cognitive representations, or constructs. Shifting between mental spaces requires moving from one cognitive representation to the other; this includes representations from different people, e.g., your own vs. another speaker's viewpoint, but also representations within a single person, e.g., expectation vs. reality or realis vs. irrealis. In an example like *If Obama would still be president, nothing like this would have happened*, the *if p then q* sets up a new mental space in which something else holds than is the case in current reality. Under such a Mental Spaces account, connectives can be treated as space-builders (Sanders, J. Sanders, & Sweetser 2009): expressions that typically establish Mental Spaces. Shifting between mental spaces has been associated with higher costs of processing a coherence relation (Segal & Duchan 1997, Segal et al. 1991, Traxler, Bybee, & Pickering 1997, Zufferey & Gygax 2015).

This section provides an overview of CCR, the discourse annotation framework used in this study, discusses the cognitive complexity of the values of each of CCR's primitives, and provides exact hypotheses regarding the marking of coherence relations based on the primitives.

4.2.1 POLARITY

POLARITY distinguishes between *positive* and *negative* relations. Positive relations are defined to hold between P and Q. Positive relations are often marked by connectives such as *also*, as in (8), *because*, or *so*. Negative relations, on the other hand, feature the negative counterpart of P or Q. Negative relations involve some kind of contrast and are often marked by connectives such as *but*, as in (9), *although*, or *however*.

- (8) [It is important that CSR, as a subject, should not be seen as a catch-all for everything that everyone wants businesses to do better.]_{S1} Also, [it is important that the core labour standards – these ILO standards – are given the dedication and focus needed for their implementation.]_{S2}
{ep-02-07-04}
- (9) [I will take the advice of the Commission on this occasion.]_{S1} but [I hope that we get it right.]_{S2} {ep-96-07-19}

Negative relations are logically more complex than positive relations, since they involve a negation on the logical operator or on one of the segments. Negative relations are also processed slower (e.g., Clark 1974, Murray 1997, Wason & Johnson-Laird 1972) and acquired later than positive relations (e.g., Bates 1976, Bloom et al. 1980, Eisenberg 1980, Evers-Vermeul & Sanders 2009). Finally, many types of negative relations involve a mental space shift, specifically relations dealing with disagreements, such as (10), in which the language user has to switch between the mental spaces of different people, or forces with conflicting outcomes, such as (11), which require a switch between two possible future worlds, and negative causal relations, in which a cause does not have its expected outcome, or vice versa, such as (12) (e.g., Pander Maat 1998, J. Sanders 1994, Verhagen 2000, 2005).

- (10) [They think you should quit.]_{S1} but [I want you to keep going.]_{S2}
- (11) [The advertisement drew a lot more people to the food truck.]_{S1} but [the kitchen could not handle the increased demand.]_{S2}
- (12) Although [Jimmy had the best voice.]_{S1} [he wasn't chosen as the band's new singer.]_{S2}

We can thus establish that negative relations are cognitively more complex than positive relations, and we hypothesize that positive coherence relations will more often be implicit than negative relations.

4.2.2 BASIC OPERATION

BASIC OPERATION refers to the operation that has to be performed to connect two discourse segments. Within this primitive, CCR traditionally distinguishes between *causal* and *additive* relations. In this chapter, we will, for simplicity's sake, employ a three-way distinction and add *conditional* relations, which were originally included as a subcategory of causal relations in Sanders et al. (1992). Both causal and conditional relations involve an implication relation, P (antecedent) → Q (consequent). In causal relations, e.g., (13), the antecedent is realized, while in conditional relations, e.g., (14), the antecedent is hypothetical. In additive relations, the segments are related to each other through logical conjunction, as in (15): P & Q.

- (13) [There is no real problem concerning this report,]_{S1} so [we can vote on it today.]_{S2} {ep-02-12-17}
- (14) [The ECOFIN Council is acting with extreme arrogance]_{S1} if [it believes it can simply disregard growth and employment in the European Union.]_{S2} {ep-96-06-18}
- (15) [Services such as tourism and clean taxes could register.]_{S1} In addition, [groups of similar small business could possibly cooperate and register together.]_{S2} {ep-00-03-14}

Since they not only involve the addition of information, but also signal an implication relation, causal and conditional relations are logically more complicated than additive relations. Conditional relations, unlike causal and additive relations, always involve a mental space shift, namely shifting to a conditional mental space (Dancygier & Sweetser 2005), and thus seem to be the most complex type of BASIC OPERATION.

The relative complexity of additive and causal relations, however, is less clear. Additive relations may be logically simpler and acquired earlier than causal relations (e.g., Bloom et al. 1980, Clark 2003, Evers-Vermeul & Sanders 2009, Katz & Brent 1968, Piaget 1924/1969), but once acquired, causal relations are processed faster (e.g., Haberlandt & Bingham 1978, Mak & Sanders 2013, Sanders & Noordman 2000) and remembered better than additive relations (Sanders & Noordman 2000, Trabasso & van den Broek 1985, van den Broek 1990). This is what Sanders (2005) has labeled the *paradox of causal complexity*, for which he proposes the causality-by-default hypothesis as a potential explanation.¹ He does point out that the expectation of a causal relation can be influenced by the characteristics of two discourse segments, so the paradox of causal complexity may be driven by only a subset of causal relations. As of yet, however, it has not been fully mapped out what features can strengthen or weaken the expectation of a causal continuation in a discourse. Crucially, the paradox

¹ The *event-indexing model* by Zwaan, Langston, and Graesser (1995) also includes something along the lines of causality being expected by default.

of causal complexity seems to hold only for positive relations, and does not pertain to relations with a negative value for POLARITY. Based on these findings, we predict no difference in the frequency with which positive additive and positive causal relations are left implicit, but we expect negative additive relations, e.g., (4) and (10), to be more often implicit than negative causal relations, e.g., (5) and (12). A summary of the exact hypotheses based on the POLARITY and BASIC OPERATION features is given in (16); positive additive and positive causal relations are expected to be most often implicit, while conditional relations are expected to be most often explicit.

- (16) positive additive, positive causal > negative additive > negative causal > conditionals

4.2.3 SOURCE OF COHERENCE

The primitive SOURCE OF COHERENCE distinguishes between relations that hold in the real world and relations that are constructed in someone's mind. CCR distinguishes two values: *objective* and *subjective* relations.² In this chapter, we adopt Sweetser's (1990) values of *content*, *epistemic*, and *speech act*. Sweetser's *content* value fully corresponds to the CCR value *objective*, but *epistemic* and *speech act* relations are grouped together under the label *subjective* in the Sanders et al. (1992) taxonomy. The main reason for adopting Sweetser's (1990) three-way distinction was the frequency with which we encountered prototypical speech act relations in the Europarl corpus³ (see also Knott 2001 for a discussion on whether to treat speech act relations as a separate category).

In a content relation, the two segments relate to each other at the propositional level. In (17), for instance, the fact that the substances cannot meet the safety standards has led to the fact that they are banned.

- (17) [These substances are banned for use in food-producing animals in the European Union]_{S1} because [it has not been possible to set a safe level of residues due to their toxic effects on human beings.]_{S2}
{ep-02-12-17}

In epistemic relations, the relation between the segments holds at the illocutionary level. Epistemic relations usually involve a speaker's reasoning process, conclusion, or judgment. In (18), for example, the first segment expresses the speaker's claim and the second segment an argument in favor of this claim.

² The SOURCE OF COHERENCE values were originally called *semantic* and *pragmatic*. These were later renamed as *objective* and *subjective*, respectively.

³ This is not entirely surprising, since the corpus consists of transcribed spoken – although in part prepared – discourse.

- (18) [Workers should be informed at the very start about the situation pertaining to their jobs]_{S1} because [they are part and parcel of that company.]_{S2} {ep-01-09-03}

In speech act relations, one of the segments relates to the speech act that makes up the other segment. The second segment in (19), for instance, motivates the asking of the question in the first segment.

- (19) [Could somebody clarify that he has actually said this please, Mr. President,]_{S1} because [it is a change of views.]_{S2} {ep-96-04-15}

Epistemic relations tend to be acquired later (e.g., Evers-Vermeul & Sanders 2011, Spooren & Sanders 2008, van Veen 2011, Zufferey 2010) and processed slower (e.g., Canestrelli, Mak, & Sanders 2013, Noordman & de Blijzer 2000, Traxler et al. 1997a) than content or speech act relations. In addition, epistemic relations usually involve a mental space shift (e.g., Canestrelli 2013, Sanders et al. 2009). In epistemic causal relations, such as argument-claim relations, there is a shift to the speaker's mental space from which the claim originates. Epistemic relations that do not involve an implication relation might not *always* involve a mental space. Examples of epistemic additive relations include two arguments in favor of the same claim or two judgments. These relations *do* involve mental space shifts if they are the first arguments or judgments listed. However, they *do not* if they are, for example, arguments or judgments numbers three and four, since the shift to the speaker's mental space took place at the claim or first judgment. So even though mental space shifting may not be all characteristic of all epistemic additive relations, it will be involved in a large number of cases.

Epistemic relations seem to be the most complex type of SOURCE OF COHERENCE, but the relative complexity of content and speech act relations is less clear-cut. There is conflicting evidence about the order in which content and speech act relations are acquired, with some studies reporting content relations being acquired before speech act relations, others speech act relations before content relations, and yet others report that there does not seem to be a difference in the age of acquisition (see Evers-Vermeul & Sanders 2011 and Zufferey 2010 for an overview). In addition, neither content relations nor speech act relations seem to involve a shift between mental spaces (Sanders, J. Sanders, & Sweetser 2009). Finally, we have not been able to find any processing studies involving speech act relations, which is probably due to their low frequency in written text (e.g., Sanders & Spooren 2015). We therefore do not predict a difference in marking between content and speech act relations.

The hypotheses regarding the implicitness of coherence relations based on the SOURCE OF COHERENCE primitive are summarized in (20).

- (20) Content, speech act > epistemic

4.2.4 ORDER OF THE SEGMENTS

ORDER OF THE SEGMENTS refers to the order of P and Q in implication relations. If the first segment conveys P and the second segment Q, as in (21), the ORDER of the relation is *basic*. If the first segment conveys Q and the second segment P, as in (22), the relation has a *non-basic* order. ORDER is not applicable to additive relations, which do not involve an implication relation and are symmetric.

- (21) Unless [we are able to do that,]_{S1} [we are going to have to pick up a huge bill.]_{S2} {ep-03-09-03}
- (22) [It has to be pursued in good times and in bad,]_{S1} unless [there is a fundamental change in the situation.]_{S2} {ep-99-01-27}

Relations with non-basic order appear to be more complex than basic order relations, since they reverse the iconic order of an implication relation. In addition, relations with iconic order seem to be easier to process (Murray 1997, Noordman & de Blijzer 2000) and acquired earlier (Evers-Vermeul & Sanders 2009) than relations involving a reversed order. We therefore hypothesize that relations with a basic order will be more often implicit than relations with non-basic order.

4.3 Parallel corpus study

The hypotheses outlined in the previous section, as well as the hypotheses concerning the influence of relational and syntactic dependency on the linguistic marking of coherence relations, have been tested using a parallel corpus study, in which we analyzed translations from English into Dutch, German, French, and Spanish. Considering multiple translation pairs makes it possible to determine whether the marking of coherence relations varies cross-linguistically. If, as we predict, the marking of coherence relations is determined by cognitive factors, we expect patterns in marking to be largely language independent.

4.3.1 Implication and implicitness

The current study uses parallel corpora to examine the marking of coherence relations. We distinguish between *implicitness*, a monolingual concept depicting something not being explicitly marked, and *implication*, a translation concept referring to something being explicitly marked in the source text (ST), but implicit in the target text (TT), the antonym being *explicitation*. In translation, there tends to be a lot of variation in the way coherence relations are expressed in the target language as compared to the source language (Halverson 2004, Zufferey & Cartoni 2014). In the field of translation studies, translators' tendency to increase the number of cohesive

ties has been hypothesized to be a universal process characterizing translated texts (Blum-Kulka 1986), but this hypothesis has only received partial support in empirical studies (Becher 2011, Zufferey & Cartoni 2014). Shifts in the marking of coherence relations, mostly explicitations but also implicitations, have also been proposed to be a by-product of the practice of sentence splitting, which may be editorially enforced (Bisiada 2016, Kruger 2017). Finally, specific instances of explicitation and implicitation in translation have also been attributed to structural or linguistic differences between languages in a language pair (Becher 2011, Fabricius-Hansen 1999, Hansen-Schirra, Neumann, & Steiner 2007).

More importantly for the research question addressed in this chapter, we assume that regardless of the exact reason for a single instance of implicitation or explicitation, changes in the marking of coherence relations should be bound by a relation's potential to be implicit. If a certain type of relation is easy to convey implicitly, source texts should contain many implicit cases of that relation, and there will be a lot of room for explicitation in translations. Similarly, that same type of relation can, if it is expressed explicitly in the source text, easily become implicit in the translation. For types of relations that are difficult to convey implicitly, on the other hand, there will be much less room for variation; if a certain type of relation is hardly ever implicit in the source text, there will be very few changes in marking between the source texts and translations. In sum, what this predicts is that the types of relations that are most often explicitated will also be the ones most often implicitated, regardless of language pair or translation direction. Conversely, the types of relations that are hardly ever explicitated will also be the ones that are hardly ever implicitated in translation. The results of a small-scale parallel corpus study reported in Hoek, Evers-Vermeul, and Sanders (2015) are in line with these predictions.

The implicitation of coherence relations in translation can thus inform us about the potential implicitness of coherence relations. Considering translations, as opposed to monolingual texts, has some advantages for the purpose of this study. First of all, it circumvents the issue of annotating implicit coherence relations. Procedures for annotating implicit relations differ between annotation frameworks and as a result, the number of implicit coherence relations identified per text can differ significantly, as becomes apparent from comparing various annotated corpora. For instance, RST tends to identify more implicit coherence relations than PDTB (compare PDTB Research Group 2007:3 and Das & Taboada 2013:9-10). The differences between annotation frameworks in the number of implicit coherence relations they identify raises the question of whether there is even a consensus on what qualifies as an implicit coherence relation. In addition, inter-annotator agreement tends to be lower on implicit relations than on explicit relations (Miltsakaki et al. 2004, PDTB Research Group 2007). Finally, using translations allows us to examine the marking of coherence relations cross-linguistically with minimal annotation effort. In the current study, we only annotated the English source text relations, while otherwise the

annotation of a substantial number of coherence relations in several languages would be required.

4.3.2 Corpus

To compare the implicitation rates of different types of coherence relations, we designed a parallel corpus using the Europarl Direct corpus (Cartoni, Zufferey, & Meyer 2013). The original Europarl corpus (Koehn 2005) consists of the proceedings of the European Parliament and their translations. Europarl Direct contains directional parallel corpora extracted from the original corpus, in which each language pair only contains speeches originally uttered in one language and the translation into another language.⁴ To be able to compare implicitation patterns across different target languages, we used four parallel corpora: EN-NL, EN-DE, EN-FR, EN-ES, of which the English ST was largely the same. We only used corpus data from before 2004 (1996-2003) to ensure that all target language fragments were direct translations; after the addition of many new countries to the European Union in 2004, the number of language pairs exceeded the translation capacity and the European Parliament started making use of pivot languages.

For the extraction of corpus fragments from the English ST, we selected eight connectives based on the type of relation they prototypically signal (see Table 1), in order to have a diverse set of relations that allows us to test our hypotheses. To be able to generalize across markers, we selected the two most frequent connectives in the corpus per relation type. We did not include *and*, since this connective is extremely general and used to signal many different kinds of relation, often as an underspecified marker (e.g., Knott & Sanders 1998, Spooren 1997).

The set of relation types considered in the current study includes many, but not all possible coherence relations. In our selection, we aimed to form a set that, based on our cognitive complexity hypothesis, was expected to cover most of the implicit-explicit spectrum (i.e., relations that would often be implicit and relations that would usually be explicit). In addition, we wanted all primitive values to be represented in as many combinations as possible in order to avoid confounds. Although this was not entirely possible (additive relations are not specified for ORDER), this was the main reason not to include temporal relations at this point; temporal relations always have positive POLARITY and can generally only have the value *content* for SOURCE OF COHERENCE.

⁴ Although members of the European Parliament are allowed to speak in any EU language, it seems that most of them, at least between 1996 and 2003, choose to speak in their native language. Out of all the English ST fragments, 85% were uttered by UK or Irish politicians. At least 50% of the other fragments were uttered by speakers who spent several years in an English-speaking country during their education. We checked all remaining fragments for grammaticality and idiomaticness.

Table 1
Selected connectives and the relations they signal

Connective (n)	POLARITY	BASIC OPERATION	SOURCE OF COHERENCE	ORDER
<i>Because</i> (388)	positive	causal	content/epistemic/SA	basic/non-b
<i>So</i> (197)	positive	causal	content/epistemic/SA	basic
<i>Also</i> (199)	positive	additive	content/epistemic/SA	N.A.
<i>In addition</i> (195)	positive	additive	content/epistemic/SA	N.A.
<i>Although</i> (254)	negative	causal/additive	content/epistemic/SA	basic/non-b/N.A.
<i>But</i> (213)	negative	causal/additive	content/epistemic/SA	basic/non-b/N.A.
<i>If</i> (254)	positive	conditional	content/epistemic/SA	basic/non-b
<i>Unless</i> (216)	negative	conditional	content/epistemic/SA	basic/non-b

We extracted all fragments that contained these connectives, along with their translations, from the four parallel corpora using Paraconc (Barlow 2008), a multilingual concordance program. Starting with the data from the English ST in one parallel corpus, we randomly selected a set of fragments, and checked whether the connectives were actually used to signal a coherence relation in each occurrence, excluding instances such as *because of*, or *because [noun]*. In addition, we removed all relations in which the connective was modified, e.g., *only because*, to keep the relations as comparable as possible, since not all of the selected connectives can be modified. We also eliminated all relations in which the connective was part of a larger fixed construction, e.g., *if only*. Finally, we removed all fragments in which *so* signaled a purpose relation, e.g., *Rick did extra homework on Wednesday so he could go to the movies on Thursday*. Andersson and Spenader (2014) demonstrate that purpose relations are rarely expressed implicitly. Since there were only six instances of purpose in our dataset, we decided to exclude them to avoid this additional confound. After finalizing our source text data set, we extracted the translations of the English fragments from the other parallel corpora. Since the corpora are not completely identical, we had to eliminate any relations for which a translation was not available in all four target languages. In the end, the selection for each connective ranged between 195 and 254 relations (see first column of Table 1), with the exception of *because*, for which we included approximately twice the number of connectives, because for other research purposes we were interested in further exploring the way in which non-basic causal relations are translated (see Chapter 5). The final dataset contained 1916 original English relations, with translations into Dutch, German, French, and Spanish.

4.3.3 Annotation of relational dependency

The discourse annotation framework used here is not the same as the annotation framework Patterson and Kehler’s (2013) model is based on (PDTB 2.0), which is why the exact annotation of the relational dependency measures cannot be replicated.⁵ In this study, we chose to determine relational dependency on the basis of explicit coherence relations only, as opposed to both implicit and explicit relations as in Patterson and Kehler’s (2013) study, to make annotations maximally comparable. We annotated relational dependency using two categories based on their impact on marking: *embedded*, for relations that are embedded in another relation, and *share/contain*, for relations that contain or share a segment with another relation.

4.3.4 Inter-annotator agreement of source text annotation

The discourse segments related to each other by the selected English ST connectives were identified, after which the relations were annotated. For the identification of the discourse segments, we used the segmentation guidelines outlined in Chapter 2. Segmentation was done prior to annotation by one annotator (author), but was open for discussion during the annotation process in case of disagreements. Per connective, we double-annotated the relations in sets of 25 (author and second expert annotator), discussing the disagreements afterwards and formulating decision rules for recurring issues. We continued this process until we reached a satisfactory agreement score on the last set ($\kappa > .7$). One annotator then finished the annotations (author). In total, 393 (20.5%) relations were double annotated using the CCR primitives.⁶ Table 2 shows the inter-annotator agreement for the annotation of the relations per primitive per relation type, of all rounds of annotation combined. Cohen’s Kappa (Cohen 1960), which is most often used to report the inter-annotator agreement in discourse annotation efforts, undervalues our agreement for certain primitives because not all categories are equally frequent. As an additional agreement measure we therefore report the AC1 metric (Gwet 2002). AC1 is comparable to Cohen’s Kappa, but is corrected for one category being more prevalent than another, i.e., asymmetric marginal probabilities. In contrast with $\text{kappa}_{\text{MAX}}$, which is also used to correct for

⁵ This difference is mostly due to implicit coherence relations. When there is no connective between adjacent sentences, PDTB has the option of adding a connective and annotating the implicit coherence relation, but also allows for other types of connections: alternative lexicalizations (AltLex), where a coherence relation “is inferred, but the insertion of an implicit connective leads to *redundancy* in its expression,” entity-based relations (EntRel), “where no discourse relation can be inferred and where the second sentence only serves to provide some further description of an entity in the first sentence,” and NoRel, “where neither a discourse relation or entity-based coherence can be inferred between the adjacent sentences” (PDTB Research Group 2007:1, original emphasis). CCR does not have these categories and would, much like for instance the RST framework, require annotating (implicit) coherence relations between most, or all, adjacent discourse segments. As a result, the number of relationally dependent coherence relations (of any kind) would be much larger when using a framework like CCR than when using PDTB.

⁶ We initially double-annotated 400 English ST relations, but seven relations were deleted from the set on the basis of the criteria outlined in Section 4.3.2.

asymmetric marginals (Feinstein & Cicchetti 1990), it does not overcorrect in case of asymmetric unbalanced marginals, which are mainly caused by annotator biases.⁷ AC1, like Cohen’s Kappa, ranges from 1 (perfect agreement) to -1 (perfect disagreement). To allow comparison of agreement scores with other discourse annotation efforts, we also include Cohen’s Kappa and the percentage of agreement.

Inter-annotator agreement was overall lowest for BASIC OPERATION within negative relations. This distinction, which roughly corresponds to distinguishing between contrastive and concessive relations, has been cited as being difficult to annotate before (e.g., Robaldo & Miltsakaki 2014, Zufferey & Degand 2017). Reaching sufficient levels of agreement on this primitive took most rounds of double annotation (five rounds in total).

Table 2
Agreement scores annotation per relation type

Relation	Primitive	AC1	Kappa	% agreement
Positive Causals	SOC	0.75	0.62	81
	ORDER	1.00	1.00	100
Positive Additives	SOC	0.73	0.62	80
Negatives	BASIC OPERATION	0.52	0.48	75
	SOC	0.75	0.62	81
	ORDER	1.00	1.00	100
Conditionals	SOC	0.67	0.52	75
	ORDER	0.88	0.88	94

In addition to the relations themselves, we annotated the relational and syntactic dependency of the English relations. Again, we double-annotated all relations in small sets until sufficient agreement was reached (author and second expert annotator). In total, all dependency measures were double-annotated for 150 relations (7.8%) marked by various connectives. The rest of the annotations were completed by one annotator (author). Table 3 shows the overall agreement scores for each dependency measure.

⁷ For a comprehensive overview of the effect of asymmetric marginal probabilities on Cohen’s Kappa and the overcorrection problem of κ_{MAX} , see Feinstein & Cicchetti (1990).

Table 3
Agreement scores annotation dependencies

Dependency	AC1	Kappa	% agreement
Syntactic Dependency	0.91	0.82	94
Relational Dependency – Embedded	0.85	0.56	89
Relational Dependency – Share/contain	0.51	0.44	73

4.3.5 Translation spotting

After annotating the English relations, we analyzed the way in which the relations were expressed in the four target languages. We categorized translations as either *explicit* or *implicit* translations based on whether or not the translated relation constituted a form of implicitation compared to the expression of the relation in the source text. We considered a relation explicitly translated if the translation included any linguistic elements that explicitly signaled a meaning equivalent to the meaning signaled by the connective in the original text. Our category of explicit translations thus extends beyond connective-to-connective translations and also includes other linguistic elements, for example prepositional phrases, verbs (e.g., causal verbs as explicit translations of causal relations), or subjunctive mood (explicitly signaling conditionality). As a result, the number of implicitations found in our study is much lower than many implicitation rates reported in NLP research, compare for instance Meyer and Webber’s (2013) 18% implicitation rate to our 4%, as presented in Table 5. For the translation spotting, we used seven labels to classify the translations, after which we categorized each translation type as either an explicit or an implicit form or translation.

Explicit connective/cue phrase, the most frequent type of translation, applies when the relation is expressed by means of a connective or cue phrase equivalent or comparable to the connective found in the source text, as in (23).

- (23) EN Mr President, **although** [it would have been better if we had voted at lunchtime,]_{S1} [I am happy to go ahead now.]_{S2} {ep-00-03-16}
- NL Mijnheer de Voorzitter, **ofschoon** we deze stemming beter hadden gehouden rond de middagpauze, vind ik het goed als we het nu doen.
- DE Herr Präsident! Es wäre **zwar** besser gewesen, wir hätten mittags abgestimmt, **aber** ich mache jetzt gern weiter.
- FR Monsieur le Président, **même** s’il aurait mieux valu que nous votions à midi, je suis ravie de voter maintenant.
- ES Señor Presidente, **aunque** habría sido mejor que hubiéramos votado a mediodía, me alegra que prosigamos ahora con la votación.

All the target languages in (23) use a connective to connect the two segments of the coherence relation. All TT connectives signal, like *although* in English, negative relations, and the resulting relations in the TT are very comparable to the original English relation.

A connective may also be paraphrased, for instance by a prepositional construction (24) or a verb (25), in which case the translations themselves do not constitute coherence relations. Other *paraphrases* maintain the coherence relation, such as the specific construction in (26), where the V1 word order, in combination with the modal verb, yields a conditional meaning in German. In all these *paraphrase* cases, the meaning of the ST connective is explicitly signaled by a linguistic device other than a connective.

- (24) EN **Unless** [we have the educational systems to teach people skills in schools and in life-long learning,]_{s1} [we will not be able to make the most of them.]_{s2} {ep-01-05-14}
FR **Sans** système éducatif efficace et sans apprentissage tout au long de la vie, nous ne pourrons en tirer pleinement parti.
'without efficient educational systems and without life-long learning...'
- (25) EN Because [it is political]_{s1} [it is more complex.]_{s2} {ep-00-02-15}
NL De politieke aard van deze uitdagingen **maakt** ze ... complex.
'The political nature of these challenges makes them complex.'
- (26) EN **If** [you are intending to postpone it to 5.30 p.m.,]_{s1} [it will totally distort the vote on this report.]_{s2} {ep-00-03-16}
DE **Sollten** Sie beabsichtigen, sie bis 17.30 Uhr zu verschieben, wird das die Abstimmung über diesen Bericht völlig verzerren.
'Were you to decide to postpone it ...'

If a relation was translated as a coherence relation, but marked by a connective that, compared to the original English connective, underspecified the relation (see for example Spooren 1997), we categorized it as *underspecified connective*. In (27), for instance, the Spanish translation of *so* is *y* "and", an additive connective. Other examples are temporal connectives signaling causal relations, causal connectives signaling conditional relations, and positive additive connectives signaling negative additive relations.

- (27) EN [The Committee of Inquiry made its point clearer than Parliament has ever been able to do,]_{s1} **so** [the public was able to read in a very clear way where the failings were.]_{s2} {ep-99-03-22}
ES La Comisión de Investigación expuso la cuestión con mayor claridad de lo que el Parlamento ha sabido hacer jamás, **y** el público pudo leer con total claridad qué era lo que iba mal.

In some instances, the segments of a relation were not related at the discourse level in the translation, but at the level of *syntax*, for instance through a relative clause construction, as in (28), or by a main clause-gerund clause construction, as in (29).

- (28) EN [The Langen report on the chemical sector is particularly important to my own constituency of Munster, and Ireland in particular,]_{S1} **because** [in Ireland at the moment there are 18,500 people employed in this sector.]_{S2} {ep-97-03-11}
- NL Het verslag-Langen over de chemische industrie is uitermate belangrijk voor mijn eigen kiesdistrict Munster en voor Ierland in het bijzonder, **waar** momenteel 18.500 mensen in de chemie emplooi vinden.
'... and Ireland in particular, where at the moment 18.500 people are employed in the chemical sector.'
- (29) EN **Because** [they were not prepared to do that,]_{S1} [they withdrew the winners of the Coca-Cola Cup, the League Cup, from UEFA's competitions.]_{S2}
- FR **N'étant** pas préparé à effectuer cette réduction, the Premier League décida alors de retirer les vainqueurs de la Coupe Coca-Cola, la Coupe de la Ligue, des compétitions de l'UEFA.
'not being prepared to implement that reduction, ...'

We labeled translations as *implicit relations* when the translation contained two discourse segments, such as in (30) that, in the absence of a connective, have to be related to each other by inference.

- (30) EN [On the one hand they claim to be controlling radioactive pollution to the marine environment, but on the other hand we have proof that there is radioactive pollution,]_{S1} **so** [there is some contradiction there.]_{S2} {ep-00-06-14}
- FR D'une part, les autorités déclarent contrôler la pollution radioactive marine, mais d'autre part, nous avons la preuve qu'il existe une pollution radioactive. Ø Nous sommes face à une contradiction.

We labeled some translations as *paraphrase-constructions*. In these translations, it is not the connective or a single segment, but the entire coherence relation that has been paraphrased. In (31), for instance, the original causal relation has been translated by an idiom made up of a prepositional verb construction. The translation does not contain a coherence relation or any marker of causality. The translation, roughly "There is no point in saying nastier things or passing nastier resolutions," does however have approximately the same meaning as the original.

- (31) EN [It is not going to change]_{S1} **because** [we say nastier things or pass nastier resolutions.]_{S1} {ep-99-01-27}
 NL Het heeft geen zin onaardige dingen te gaan zeggen of onaardige resoluties te gaan aannemen.

A part of the translations was double-coded to make sure the annotations were reliable, using the same step-wise process as described in Section 4.3.4 for the annotation of the source text. The inter-annotator agreement of the translation spotting per target language is given in Table 4.

Table 4
Agreement scores translation spotting per target language

Target language	AC1	Kappa	% agreement
Dutch	0.94	0.75	94
German	0.96	0.82	96
French	0.96	0.73	96
Spanish	0.95	0.36	96

The seven types of translations we distinguished while translation spotting were collapsed into three categories for the final analysis: *explicit translations*, *implicit translations*, and *other* (see Table 5 for an overview). We considered TT fragments labeled as *explicit connective/cue phrase* or *paraphrase* to be instances of *explicit translations*, since both translations types involve an explicit linguistic expression indicating the type of relation that holds between the segments.

We grouped translations categorized as *underspecified connective* with either the implicit or explicit translations depending on which factors we considered in our analysis. There is a crucial difference between the predictions made by the cognitive complexity hypothesis as outlined in Section 4.2, and the hypotheses of the effects of the different dependency measures, presented in Section 4.1.2. While the dependency measures predict the likelihood of the presence of a connective, the hypotheses formulated on the basis of the cognitive complexity of relational features concern whether or not a specific feature is explicitly signaled. In (27), for example, the Spanish translation is explicit in the sense that it contains a connective. The causality, however, is implicit and has to be inferred. When taking into account only the type of coherence relation, underspecified connectives were therefore grouped with the implicit relations. When also taking into account the dependency measures, underspecified connectives were considered to be explicit translations.

We considered both *implicit relation* and *syntax* cases to be instances of *implicit translations*. Although the translations in the *syntax* category do not contain

a prototypical implicit coherence relation, they do contain two clauses between which the type of relation is not explicitly signaled by means of a connective.

Two types of translations could not be assigned to either label and were left out of the final analyses: *paraphrase-construction* and *other*. Since in *paraphrase-constructions* the entire coherence has been paraphrased, it is impossible to categorize them as either implicitations or explicitations. Translations of this type were therefore left out of the final analyses. Translations labeled as *other* were fragments in which the meaning of the text was changed dramatically, for instance if a totally different type of connective has been used to mark the relation, e.g., a negative connective instead of a positive conditional connective, or if one of the segments had not been translated, or was missing from the corpus.

4.4 Results

We used binary logistic regression, a method for modelling a dichotomous categorical outcome variable using a set of predictor variables, to model the *implicit* versus *explicit* marking of coherence relations, using R (R Core Team 2016, version 3.2.4). Table 5 contains an overview of the frequency of the translation categories in the entire dataset.

Table 5
Types of explicit and implicit translations and their frequency

Type of translation	Explicit/Implicit	Frequency (%)
Explicit connective/cue phrase	Explicit	91.5
Paraphrase	Explicit	2.3
Underspecified connective	Explicit / Implicit	0.5
Syntax	Implicit	0.3
Implicit relation	Implicit	3.2
Paraphrase-construction	N.A.	1.6
Other	N.A.	0.6

As explained in Section 4.4.3.5, the categorization of *underspecified connectives* as either implicit or explicit is different when taking into account only predictions made on the basis of the CCR primitives (implicit) than when considering the influence of relational and syntactic dependency on the marking of coherence relations (explicit).

We first built a model using only target language and the four CCR primitives as predictor variables, grouping underspecified connectives with *implicit*. We used a combination of forward and backward step-wise selection to arrive at the best model. The final model was selected using the Akaike Information Criterion (AIC). We performed post-hoc tests using the multcomp package (Hothorn, Bretz, & Westfall

2008) to obtain additional pairwise comparisons for all variables with more than two levels.

The final model included *target language*, *POLARITY*, *BASIC OPERATION*, *SOURCE OF COHERENCE*, and *ORDER*, and an interaction of *POLARITY* and *BASIC OPERATION*.⁸ The main model is presented in Table 6. The model includes an intercept for each target language, which represents the explicit translations for the reference categories of each primitive, i.e., for positive content causal relations with basic order. A positive regression coefficient is to be interpreted as meaning that the category increases the odds of a relation being explicitly marked; a negative regression coefficient is to be interpreted as meaning that the category increases the odds of a relation being implicit. For instance: a relation has a higher chance of being explicit when it has a negative value for *POLARITY* than when it has positive *POLARITY* ($\beta=1.24$, $SE=0.25$, $p<.001$). Note that in the model, the data have been linearized by means of a logit transformation, which means that while the influence of a certain parameter on *marking* may be the same for all target languages in terms of logits, the increase in logits does not affect the *proportion* of explicit translations into each target language in the same way.⁹ As an example: a negative value for *POLARITY* increases the chances of a relation being explicit by 1.24 logits. For Dutch, this means that the proportion of explicit translations increases from 91.91% to 97.4%, but for French, the proportion of explicit translations increases from 93.88% to 98.15%.

The overall number of implicitations varied between target languages, with the least implicitations from English into Spanish; there were fewer implicitations into Spanish than into Dutch ($z=7.31$, $p<.001$), German ($z=5.34$, $p<.001$), or French ($z=5.86$, $p<.001$). Crucially, however, we found no interaction between *target language* and any of the other predictor variables.

In line with our predictions, coherence relations with basic order became implicit more often than relations with non-basic order. Speech act relations were more often implicated than content relations. They were also more often implicated than epistemic relations ($z=-2.96$, $p<.01$). There was no difference in implication between content and epistemic relations. This is only partly in line with our hypothesis, since we predicted both content and speech act relations to appear implicit more often than epistemic relations.

⁸ There was one competing model, which included a three-way interaction of *POLARITY*, *BASIC OPERATION*, and *SOURCE OF COHERENCE*. However, since the majority of the extra parameters in this model were otiose, we selected the more parsimonious model.

⁹ The proportion of explicit cases can be calculated using the following formula: $\frac{1}{1+e^{-\beta}}$

Table 6
Regression coefficients and their standard errors and p-values

	Parameters	β	SE	<i>p</i>
<i>Target language</i>	Dutch	2.43	0.16	<.001
	German	2.83	0.17	<.001
	French	2.73	0.16	<.001
	Spanish	4.10	0.24	<.001
<i>POLARITY BASIC OPERATION</i>	Δ Negative	1.24	0.25	<.001
	Δ Additive	0.20	0.17	.23
	Δ Conditional	3.25	0.51	<.001
<i>POLARITY* BASIC OPERATION</i>	Δ Positive additive	0.20	0.17	.23
	Δ Positive conditional	3.25	0.51	<.001
	Δ Negative causal	1.24	0.25	<.001
	Δ Negative additive	0.48	0.20	.02
	Δ Negative conditional	3.16	0.59	<.001
<i>SOURCE OF COHERENCE</i>	Δ Epistemic	0.03	0.13	.81
	Δ Speech act	-0.53	0.20	<0.01
<i>ORDER</i>	Δ Non-basic	-0.75	0.16	<.001

Table 7
Additional pairwise comparisons for the interaction of *POLARITY* and *BASIC OPERATION*

Parameter	Reference category	<i>z</i> -score	<i>p</i>
Positive conditional	Positive additive	5.73	<.001
Negative causal	Positive additive	3.59	<.01
Negative additive	Positive additive	1.38	.70
Negative conditional	Positive additive	4.92	<.001
Negative causal	Positive conditional	-3.66	<.01
Negative additive	Positive conditional	-5.09	<.001
Negative conditional	Positive conditional	-0.11	1.00
Negative additive	Negative causal	-2.45	.12
Negative conditional	Negative causal	3.09	.02
Negative conditional	Negative additive	4.40	<.001

The model includes an interaction between *POLARITY* and *BASIC OPERATION* that is largely in line with the hypothesis formulated in (16) in Section 4.2.2. The additional pairwise comparisons between the levels of this interaction are given in Table 7. Both positive (e.g., (14)) and negative conditional, (e.g., (21) or (22)) relations were more often explicit than the other types of relations; there was no difference in the frequency with which positive and negative conditional relations were implicitated. Positive causal relations (e.g., (13) or (18)) became implicit more often than both negative causal (e.g., (5) or (12)) and negative additive relations (e.g., (4) or (10)). Positive additive relations (e.g., (8) or (15)) were more often implicit than negative causal relations. There was neither a difference in implicitation between positive causal and positive additive relations, nor between negative causal and negative additive relations.

We also built a model for which we considered target language, the CCR primitives, and the syntactic and relational dependency measures. For this analysis, we grouped the underspecified connectives with the explicit cases. The best model was very similar to the model in Table 8 and included *target language*, *SOURCE OF COHERENCE*, *ORDER*, an interaction of *POLARITY* and *BASIC OPERATION*, and *share/contain*. As hypothesized, relations that were relationally dependent in the sense that they contained another relation or shared a segment with another relation became implicit more often than relations that were not ($\beta=-0.41$, $SE=0.13$, $p<.01$).

The final model did not include *syntactic dependency* or *embedded* parameters. In an otherwise empty model, both were significant in the direction that we hypothesized, with dependent relations being implicit less often. These effects disappeared after we added other parameters, and leaving out *syntactic dependency* and *embedded* resulted in a better model than leaving out other variables. On closer inspection, it appears there is a strong relation between both these dependency measures and *BASIC OPERATION*. The type of *BASIC OPERATION* that is least often implicit (conditional) is also the type of relation that is most often embedded under a syntactic construction ($\chi^2(2)=163.83$, $p<.001$), as in (32), or in another coherence relation ($\chi^2(2)=207.73$, $p<.001$), as in (33), see Table 8. Conversely, the types of *BASIC OPERATION* that are most often implicit – causal and additive relations – also most often share a segment with, as in (34), or contain another relation, as in (35), ($\chi^2(2)=148.04$, $p<.001$).

- (32) We believe that **if** [our aim is to globalize not only the market in all kinds of goods,]_{s1} [we should encourage international work-sharing to the profit of poorer economies and ourselves.]_{s2} {ep-97-11-21}
- (33) Mr President, we can write all we like about human rights and we can report and discuss the issue, but **unless** [we follow through on our intentions and the reports,]_{s1} [they are meaningless.]_{s2} {ep-00-03-15}

- (34) [We will come back again with material on these issues,]_{S1} **because** [these problems, in one way or another, are here to stay,]_{S2} simply because development cooperation represents basically doing something in the world that is different from taking care of traditionally-defined trade interests, etc. {ep-00-02-16}
- (35) [The scope of application should be limited to serious offences linked to organised crime.]_{S1} **In addition**, [it is not the relationship between the convicted person and the third party which is important, but the fact that illegally obtained property was transferred to a third person.]_{S2} {ep-02-11-18}

Table 8

Number of dependent and independent relations, per BASIC OPERATION, per type of dependency

	<i>Syntactic dependency</i>		<i>Relational dependency: Embedded</i>		<i>Relational dependency: Share/contain</i>	
	Dependent	Independent	Dependent	Independent	Dependent	Independent
Causal	86	668	127	627	324	430
Additive	52	634	89	597	371	315
Conditional	94	376	138	332	87	383

4.5 Discussion and conclusion

This chapter set out to investigate which factors make coherence relations more, or less, likely to remain implicit. Starting from the assumption that cognitively simple relations are more often implicit than relations that are cognitively more complex, we used the primitives of the Cognitive approach to Coherence Relations to formulate fine-grained hypotheses about which types of coherence relations would be most often implicit, and which relations would often be explicitly marked. These hypotheses were tested by means of a parallel corpus study. The results seem to be in line with the idea that the linguistic marking of coherence relations is influenced by the cognitive complexity of the relation, as determined on the basis of its values for each of the four CCR primitives using evidence from logic, language acquisition, language processing, and Mental Space theory.

Specifically, we found that coherence relations with basic order are more often implicit than relations with non-basic order and that conditional relations are less often implicit than causal or additive relations. In addition, we found that the types of coherence relations that are often implicitated did not vary between the languages in our corpus. Since complexity is a general cognitive concept, we had indeed expected our results to hold cross-linguistically.

One of the findings that was only partly in line with our hypotheses was that speech act relations were more often implicit than both content and epistemic

relations. A potential explanation for speech act relations frequently being implicit could be that they are easy to recognize, since one of the segments often involves a question or an imperative. As such, the number of possible relations that can be constructed between the speech act and the preceding or following segment may be limited, making additional signaling redundant. Furthermore, there are arguments in favor of the relative simplicity of speech act relations: they stay close to the I-here-now of the communicative situation in which they are expressed, and are always directly connected to a speaker or author (Sanders et al. 2009). However, speech act relations are a rarely investigated type of coherence relation, which makes it difficult to come up with a conclusive explanation at this point.

As hypothesized, we did not find that positive causal relations are more often implicit than positive additive relations, as would be expected if people actually *by default* expected causal relations of any type. That does not seem to be the case, as was already suggested in Sanders' (2005) proposal of the causality-by-default hypothesis. Instead, it appears that people's causal expectations may be more specific, potentially driven by contextual features, and possibly limited to a certain type of causal relations. When people expect causal relations and what they base these expectations on seems a fruitful future research endeavor.

The current study also included relational and syntactic dependency as additional measures that may affect the implicit versus explicit marking of coherence relations. The hypotheses formulated with regard to these dependency measures were only in part confirmed; while we did find that relations that shared a segment with another relation (*share*) and relations that contained another relation within one of their segments (*contain*) would more often become implicit, we did not find that relations that were either embedded in another relation (*embedded*) or in a syntactic construction (*syntactic dependency*) were more often explicit. Our study indicates that the dependency measures are not completely independent of, or orthogonal to *relation type*. Crucially, the types of relations that are most often implicit –causal and additive relations– are also the types of relations that most often share a segment with another relation or contain another relation in one of their segments. Conversely, the types of relations that are least often implicit –conditional relations– are also the types of relations most often embedded in another relation or a syntactic construction. This can explain why *share/contain* is a significant parameter in our model, and *embedded* and *syntactic dependency* are not. While there is a lot of variation in marking to be explained for *share/contain*, this is not the case for *syntactic dependency* and *embed* since the vast majority of those relations is explicit anyway.

The relational dependency measures used in this chapter were based on Patterson and Kehler (2013), who do report a significant contribution of all relational dependency types. There are a few possible explanations for why their results differ from the results reported in this chapter. First of all, the set of relation types under investigation overlap, but are not the same. Specifically, Patterson and Kehler's study

does not include conditional relations. In addition, they use monolingual data, in which the percentage of implicit relations is higher than the number of implicitations in this study (see PDTB Research Group, 2007:65-90), and which therefore contains more variation to be explained. Finally, Patterson and Kehler treat the relational dependency measures as a single group in their model. While the group as a whole may be a significant predictor, the contribution of the individual parameters is not clear.

The current study has demonstrated that there seems to be a certain link between contextual dependency, both relational and syntactic, and relation type. This suggests that the way in which different types of coherence relations tend to be embedded in the rest of the discourse is in some way structured. The exact nature of the connection between contextual dependency and relation type would be worthwhile exploring in the future.

This study used parallel corpora to investigate the implicit versus explicit expression of coherence relations. As was explained in Section 4.3.1, a relation's often becoming implicit in translation is grounds to assume that this type of relation is more often implicit in monolingual data as well. Parallel corpora can therefore not just inform us about translations, but serve as a valuable resource for investigating monolingual phenomena as well. The main benefit of our study is that it allows us to generalize across languages and to explore whether the features influencing the marking of coherence relations are cross-linguistic or language specific. There are, however, also a few drawbacks to this approach. First of all, the current study does not include implicit-to-implicit translations. This means that types of relations that are hardly ever expressed explicitly might not be included in our corpus. Future studies that further explore the cross-linguistic expression of coherence relations might want to make sure they include these types of relations as well. In addition, based on this study we are only able to say something about the relative differences in marking, since our findings are based on translation data. Absolute differences in the frequency with which relations are explicitly marked would require the annotation of comparable monolingual data, in several different languages. Furthermore, we used the Europarl corpus, which is a fairly specific genre that includes a mix of formal and informal discourse and involves a hybrid form of speech and writing, since it is a transcription of spoken but partly prepared discourse. Taboada (2006) proposes that the marking of coherence relations may in part be dependent on genre and this finding was corroborated by Dupont and Zufferey (2017). We have as of yet not been able to find or compile another translation corpus on which we can replicate this study. However, in as far as the data are comparable, the patterns in the marking of coherence relations found in the current study are very similar to patterns reported on the basis of the PDTB (Asr & Demberg 2012) and the RST Treebank (Taboada & Das 2013), both of which use newspaper texts from the Wall Street Journal.

In this chapter, we have proposed that the linguistic marking of coherence relations is influenced by cognitive complexity. In addition, our study is in line with the idea that contextual dependency also affects whether a relation is explicitly marked. It is, however, very likely that other factors also play a role. Certain elements in the segments or context of a relation may for instance function as indicators of a specific relation, as is most notably explored by Das, Taboada, & McFetridge (2015). Conversely, some elements in the segments or context of a relation may block other, more default interpretations. All of these elements, and potentially more, are crucial in coming to a comprehensive overview of the complicated issue of when, how, and why we mark coherence relations.

5 How connectives and segment-internal elements interact in the linguistic marking of coherence relations

Connectives and cue phrases are the most prototypical linguistic elements that signal coherence relations, but by limiting our attention to connectives, we are likely missing out on important other cues readers and listeners use when establishing coherence relations. However, defining the role of other types of linguistic elements in the signaling of coherence relations is not straightforward, and it is also not obvious why and how non-connective elements function as signals for coherence relations. In this chapter, we aim to develop a systematic way of categorizing segment-internal elements as signals of coherence relations on the basis of a literature review and evidence from parallel corpora. We propose a three-way distinction between division of labor, agreement, and general collocation to categorize the different ways in which elements inside discourse segments interact with connectives in the marking of coherence relations. In each type of interaction, segment-internal elements can function as signals for coherence relations, but the mechanism behind it is slightly different for each type.

This chapter has been submitted as:

Hoek, Jet, Zufferey, Sandrine, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (submitted). *The linguistic marking of coherence relations: Interactions between connectives and segment-internal elements.*

5.1 Introduction

When readers or listeners are presented with a text, they do not treat the individual clauses and sentences in that text as independent and unrelated. Instead, they try to relate each part of the text, or each *discourse segment*, to the rest of the discourse. By doing this they, ideally, create a coherent mental representation of the discourse. Discourse segments can be related to each other in different ways, for instance through a causal relation, a contrastive relation, or a conditional relation, so language users have to figure out the exact way in which a segment has to be related to another segment. The process of arriving at an appropriate type of coherence relation is often, but not always, facilitated by the presence of explicit linguistic cues in the discourse. The most obvious markers of coherence relations are discourse connectives (e.g., *before, if*) and cue phrases (e.g., *for this reason, by contrast*). Coherence relations that are marked by a connective or a cue phrase, as in (1a), are commonly referred to as *explicit* coherence relations; coherence relations that are not accompanied by a connective or a cue phrase, as in (1b) are traditionally labeled *implicit* coherence relations. Even though (1b) lacks the explicit instruction to establish a causal relation that *because* provides in (1a), the relation between the two discourse segments is still easily inferable.

- (1a) [Trey pushed Tara.]_{S1} **because** [she threw his baseball bat into the water.]_{S2}
 (1b) [Trey pushed Tara.]_{S1} \emptyset [She threw his baseball bat into the water.]_{S2}

Although the distinction between explicit and implicit coherence relations seems very straightforward, it is not without its problems. Connectives can for instance signal a relation that is less specific than the relation that is constructed by language users. (1c), for example, is marked by the temporal connective *after*, but it is very likely that a causal relation will still be inferred between its segments. While both (1a) and (1c) are marked by a connective, they thus seem to differ in the extent to which the relation is explicitly signaled.

- (1c) [Trey pushed Tara.]_{S1} **after** [she threw his baseball bat into the water.]_{S2}

Similarly, relations that do not contain a connective or a cue phrase can vary in the extent to which the relation remains linguistically implicit. The relation in (2), like the relation in (1b), does not contain a connective, but (2) seems to be crucially different from (1b) in that the relation contains other linguistic cues that help in constructing the relation: the semantic opposites *great* and *horrible*. (2) therefore appears to be less implicit than (1b).

- (2) [Jack is a great kid.]_{S1} Ø [His sister is horrible.]_{S2}

When linguistic cues other than connectives or cue phrases are taken into account, the distinction between implicit and explicit relations becomes more fuzzy, and the category of implicit coherence relations becomes much smaller.

5.1.1 The marking of coherence relations by connectives

Connectives and cue phrases (from now on referred to as ‘connectives’) provide processing instructions about the way in which two discourse segments should be related to each other (e.g., Sanders & Spooren 2007). The general principle behind the marking of coherence relations seems to be that if the appropriate relation can be constructed without a connective or with an underspecified connective, it can be left implicit or underspecified; if not, the relation should be explicitly marked. This idea is compatible with several different theories of language production. Grice’s (1975) Quantity maxim, for instance, states that a contribution should neither say more nor less than necessary. This two-part maxim has been reformulated by Horn (1984:13) in terms of speaker-based and hearer-based principles: the Q principle, which states that you have to “make your contribution sufficient,” and the R principle, which states that you should “make your contribution necessary.” While the Q principle reduces the hearer’s effort to interpret an utterance, the R principle prevents the speaker from wasting effort producing superfluous linguistic content. A coherence relation should therefore be sufficiently marked so that the hearer will be able to construct the appropriate relation, but not be overly or unnecessarily marked so as to limit the speaker’s efforts.

The same idea can also be thought of in terms of in effort versus effect, key notions from Relevance Theory (Sperber & Wilson 1985, Wilson & Sperber 2005); if inferring a specific type of relation requires too much effort for the resulting cognitive effect, another, easier relation will be inferred. Explicitly marking a relation that would otherwise be misinterpreted ensures that the right relation is constructed by the hearer. From a speaker perspective, adding a connective to a relation that could already be inferred without it costing effort, but the resulting effect is minimal.

Finally, the Uniform Information Density (UID) hypothesis (Frank & Jaeger 2008, Levy & Jaeger 2007) proposes that speakers “structure their utterances so as to avoid peaks or troughs in information density” (Levy & Jaeger 2007:1). If a coherence relation can be established without the presence of a connective or cue phrase, inserting a connective or cue phrase would lead to a trough in information density, since it does not add any extra information to the utterance. Conversely, leaving a relation that is hard to interpret without a connective results in a peak in information density, since too much information has to be extracted from the other linguistic elements in the utterance.

All these accounts predict that speakers use a connective when it contributes essential information to the discourse. If the connective is barely informative or even entirely redundant, speakers will be more inclined to leave it out, in which case the relation will be implicit, or use a more general connective, in which case the relation will be underspecified. If another element within a discourse segment already signals or partly signals how that segment should be related to another segment from the discourse, this would eliminate or reduce the amount of information a connective would contribute. Other linguistic elements that convey information or raise expectations about the type of coherence relation that should be constructed are thus expected to influence the marking of coherence relations by connectives.

5.1.2 The marking of coherence relations by other linguistic elements

Research on the marking of coherence relations has mostly been focused on connectives and cue phrases, because these are the only linguistic elements that by definition express relational meaning; they signal the meaning that the relation adds to the meaning of the segments in isolation. As such, connectives and cue phrases have a special status in cognitive and linguistic studies of coherence relations, especially because there are systematic restrictions on their meaning and use and because these systems organize the lexica of connectives in various languages (e.g., Knott & Dale 1994, Knott & Sanders 1998). Connectives and cue phrases being the only elements that exclusively express relational meaning does not necessarily mean, however, that connectives and cue phrases are the *only* indicators for coherence relations. By limiting our attention to connectives, we are therefore likely missing out on important other cues readers and listeners use when establishing coherence relations.

The most elaborate research effort to identify other signals for coherence relations has been the recently released RST Signalling Corpus (Das, Taboada, & McFetridge 2015), in which linguistic cues that signal coherence relations annotated in the RST Treebank (Carlson, Okurowski, & Marcu 2002) are identified. However, since linguistic elements other than connectives do not directly signal coherence relations, it is not obvious how to group or categorize signals. For example, Taboada and Das (2013:258) indicate that the numerical element *five* in the first segment of (3) is a signal for the ELABORATION relation between S_1 and S_2 , with the five names in S_2 specifying the contents of the “five.”¹ The relation in (3) is then categorized as being signaled by a *numerical* signal.

¹ It could, however, be argued that the set of five in S_1 is a different set than the five names in S_2 .

- (3) [This maker of electronic devices said it replaced all **five** incumbent directors at a special meeting ...]_{S1} [Elected as directors were Mr. Hollander, Frederick Ezekiel, Frederick Ross, Arthur B. Crozier and Rose Pothier.]_{S2}

Although it seems indeed plausible that *five* in S₁ and the list of five names in S₂ are in this case important linguistic elements when inferring a coherence relation between these two segments, the *numerical* signal does not in itself signal an ELABORATION relation. *Five*, or any other number, may just as well occur in any other type of relation. In (4), the number *five* is for instance in contrast with the number *twelve*, and could as such be taken to signal a CONTRAST relation.

- (4) [Jane packed five pairs of socks for her school trip.]_{S1} [Frank brought twelve pairs.]_{S2}

Instead of considering both (3) and (4) as instances of *numerical* signaling, they could receive more specific labels. (3) could for instance be categorized as *numerical-lexical enumeration* and (4) as a *numerical pair*, which would allow us to distinguish between these two relations, while at the same time attributing signal status to other parts of the relation as well (the list of names and *twelve*, respectively). However, other examples may provide new ideas for labels and groupings. In sum, determining which parts of a relation function as a signal is not straightforward, nor is it obvious how to categorize the identified signals.

While the RST Signalling Corpus is a valuable inventory of potential signals, both within the segments and outside them (e.g., text structure, genre), it does not link signals to relation types in a systematic way and does not comment on *how* or *why* the indicated signals function as cues for coherence relations. In addition, since the annotation was mostly focused on relations without connectives, the RST Signalling Corpus does not identify potential additional signals in relations that contain a connective. The current chapter explores the marking of coherence relations by connectives on the one hand, and other types of cues on the other. Specifically, we will investigate how linguistic elements within the segments of a coherence relation, i.e., *segment-internal elements*, can contribute to the marking of the relation, and how the presence of other signals relates to the presence of connectives.

Connectives encode instructions on how to relate discourse segments to each other, which is why they themselves are not part of the segments. Throughout this chapter, connectives will therefore be left out of the segmentation of the relations whenever possible. Often, the marker will appear either before S₁ or between S₁ and S₂ (i.e., at the head of S₂), but this is by no means a categorical feature of connectives, since they can also appear in clause-medial or clause-final position (in S₂). The class of connectives includes several different grammatical categories, such as

conjunctions, adverbs, and idiom chunks. While conjunctions are restricted to appear either between the two segments or before S_1 , adverbs and idiomatic phrases can also appear in clause-medial or clause-final position. This is a crucial distinction, since even if connectives appear in clause-medial position, this does not make them part of the segment, or segment-internal cues. We consider elements to be segment internal if they are integrated in and are part of the propositional content of the clauses that are, or are part of, the segments of a coherence relation.

5.1.3 The interaction between connectives and segment-internal elements

There are several segment-internal features that have been linked to particular types of coherence relations. These segment-specific elements include a wide range of linguistic categories, such as complex phrases, lexical items, modal markers, and verbal inflection. The features can either occur in one of the segments or in both of the segments. However, it seems that not all linguistic elements that have been associated with a specific type of coherence relation can function as a signal for the relation in the same way, and there appear to be differences in the way in which the presence of a specific linguistic element in one of the segments of a relation can impact the marking of that relation by means of a connective. In this chapter, we argue that there are three distinct ways in which segment-internal elements systematically interact with the connective that marks a coherence relation, which we label *division of labor*, *agreement*, and *general collocation*.² In all three types of interactions, the segment-internal element can function as a cue for which type of coherence relation should be constructed or expected, but the reason why the segment-internal element functions as a signal for a coherence relation is slightly different in each interaction. The three-way categorization is based on specific linguistic elements that have been reported to often co-occur with specific types of connectives or coherence relations in monolingual corpus studies or experimental studies; the existence of the three types of interactions is demonstrated using data from parallel corpus data. Section 5.2 first introduces the parallel corpus used in this study and the rationale behind using translations to study the marking of coherence relations. Sections 5.3, 5.4, and 5.5 discuss, respectively, *division of labor*, *agreement*, and *general collocation*.

5.2 The marking of coherence relations in parallel corpora

The current chapter makes use of parallel corpus data to investigate the marking of coherence relations by connectives and cues phrases and segment-internal elements. Section 5.2.1 explains how translations can provide insight into which parts of a text fragment contribute to the marking of a coherence relations and introduces the specific

² Technically speaking, division of labor and agreement interactions could also be considered examples of collocations, either between a segment-internal element and a relation type (as in division of labor) or between a segment-internal element and a specific connective and relation type (as in agreements).

discourse-annotated parallel corpus used in this chapter. Section 5.2.2 gives a brief overview of the annotation model used to annotate the coherence relations in the parallel corpus.

5.2.1 Coherence relations in translation

Monolingual corpora are extremely valuable resources for language research, but when studying *meaning* using a monolingual corpus, researchers still mainly have to rely on their own interpretations, since “meaning is not directly observable,” (Noël 2003:758). When it comes to the interaction between segment-internal elements and connectives, it is not necessarily obvious what and how each element contributes to the overall interpretation of a relation. A proposal for an alternative method to research meaning is to make use of parallel corpora, which consist of a source text (ST) and one or multiple translations (target texts: TTs) (e.g., Dyvik 1998, Melamed 2001, Noël 2003, Teubert 1999). In this approach, the translator is treated as a naive ‘annotator,’ whose main purpose was to accurately convey the meaning of the ST in the TT. Variety in the linguistic means used to arrive at a similar meaning can inform researchers about the contribution of individual elements or constructions to the meaning of a text fragment.

Connectives are known to be volatile items in translation (Halverson 2004, Zufferey & Cartoni 2014). Translations often use an equivalent connective to translate the original connective, as in (3) but can also leave out the connective entirely, as in (4), rephrase the meaning of the connective using different linguistic means, as in (5), or use a connective that does not express the same meaning as the original connective, as in (6).

- | | | |
|-----|----|--|
| (3) | EN | [This is nothing to do with aeroplanes,] _{S1} although [I could speak about flights if you wished me to.] _{S2} |
| | ES | [Esto no tiene nada que ver con aviones,] _{S1} aunque , [si lo desean, puedo hablar de volar.] _{S2} |
| (4) | EN | [I would like a clear answer from the Commission as to why it is failing to do what this Parliament has asked it to do and what it was instructed to do by the Council at the spring summit,] _{S1} because [this is not acceptable.] _{S2} {ep-02-11-20} |
| | DE | [Ich bitte die Kommission um eine klare Antwort, warum sie der Forderung des Parlaments und der Anweisung des Rates, die ihr auf dem Gipfeltreffen im Frühjahr erteilt wurde, nicht nachkommt.] _{S1}
Ø [Dies ist nicht hinnehmbar.] _{S2} |

- (5) EN [Currently there are no equivalent tests,]_{S1} **so** [those tests are for the moment standard.]_{S2} {ep-02-06-13}
- FR [Il n'existe pas à l'heure actuelle de tests équivalents,]_{S1} [**ce qui leur confère la qualité** de tests standards.]_{S2}
'Currently there are no equivalent tests, which gives them the status of standard tests.'
- (6) EN [I think honourable Members are aware at least of the organizational complexity]_{S1} **because** [all this has major practical and legal implications.]_{S2} {ep-96-11-19}
- NL [Dit brengt praktisch en juridisch gezien heel wat met zich mee,]_{S1} **en** [ik denk dat de geachte afgevaardigden wel inzien hoe moeilijk dit te organiseren is.]_{S2}
'This has major practical and legal implications, and I think the honourable Members are aware of how difficult this is to organize.'

The translations in (3)-(6) show how coherence relations can vary in the extent to which they are explicitly signaled, with the translated relations in (4) and (6) being less explicit than the English originals, and the French translation in (5) being arguably even more explicit than the original English relation. The variations introduced in the linguistic encoding of the marking of coherence relations in translation supposedly have no detrimental effects on the meaning of a text as a whole. This makes translation data especially suited to research the different linguistic resources that can be used to signal coherence relations.

In the current chapter, we make use of a multilingual discourse-annotated parallel corpus (see Chapter 3) based on the Europarl Direct corpus (Cartoni, Zufferey, & Meyer 2013, Koehn 2005) to supply additional evidence for the existence of three different types of interactions between connectives and segment-internal elements. The corpus consists of approximately 2000 explicit English coherence relations (with the connectives *also*, *although*, *because*, *but*, *if*, *in addition*, *so*, *unless*) with translations into Dutch, German, French, and Spanish. All English relations were annotated using CCR (Cognitive approach to Coherence Relations, see Section 5.2.2). We hypothesize that the different types of interactions between segment-specific elements and connectives can be distinguished in translation and formulate specific expectations for the translation of each type of interaction in the three corpus data sections (Sections 5.3.2, 5.4.2, and 5.5.2).

5.2.2 The Cognitive approach to Coherence Relations

The coherence relations in the English part of the parallel corpus used in the chapter have all been annotated using the using the cognitive approach of coherence relations (CCR; Sanders, Spooren, & Noordman 1992 and later work; see also Chapter 3). This

framework uses a set of primitives, rather than single end labels to classify coherence relations. Each basic primitive has at least two values: POLARITY (positive, negative), BASIC OPERATION (causal, conditional, additive), SOURCE OF COHERENCE (objective, subjective), ORDER OF THE SEGMENTS (basic, non-basic), and TEMPORALITY (temporal, non-temporal) (Evers-Vermeul, Hoek, & Scholman 2017, Sanders et al. 1992).³ In addition to the basic primitives, several other features have been identified that, unlike the four main primitives, apply only to a small subset of relations, e.g., VOLITIONALITY and PURPOSE (see Chapter 3).

CCR and its primitives were originally developed to depict coherence relations. They can, however, also be used to describe what is signaled by connectives (e.g., Knott & Sanders 1998). Connectives provide language users with instructions on how to relate two (or, occasionally, multiple) discourse segments to each other, but not all connectives are equally specific. *Because*, for instance, signals a positive value for POLARITY and a causal value for BASIC OPERATION, while the only thing that *but* signals is a negative value for POLARITY. As the relation in (1c), repeated below as (7), illustrates, it is possible for a connective to appear in a coherence relation that is more specific than what it in itself signals; *after* signals positive temporal relations, but the coherence relation in (7), i.e., the relation in the mental representation constructed on the basis of (7), is more appropriately labeled a positive causal relation.⁴

(7) [Trey pushed Tara]_{S1} **after** [she threw his baseball bat into the water.]_{S2}

Since CCR can be used to describe what is signaled by connectives, it should also be applicable to depicting what is signaled by other linguistic elements that contribute to the marking of coherence relations. The semantic opposition in (2), for instance, indicates that the relation has a negative value for POLARITY.

In the remainder of the chapter, the CCR primitives will be used to depict the type of coherence relations that hold between two discourse segments, as well as to describe what is explicitly signaled by connectives and other signals of coherence relations. The individual primitives allow us to determine to what extent linguistic cues, connectives or other elements, explicitly signal the relation.

5.3 Division of labor

There are several segment-specific cues that appear in one or both of the segments of a coherence relation that seem to eliminate or reduce the need for a connective because

³ See Sanders et al. (1992) for the original CCR taxonomy, and Chapter 3 for a more elaborate overview of the considerations and decisions made while annotating the English ST relations in the dataset used in the current study.

⁴ Causal relations are more specific than temporal relations, see e.g., Evers-Vermeul, Hoek, and Scholman (2017), Sanders et al. (1992), Spooren (1997).

the segment-internal elements themselves already signal which type of relation should be inferred the two discourse segments. In these cases, there appears to be a *division of labor* between the segment-internal cue and the connective; either the connective signals the relation, or an element inside one or both of the segments does so. It even seems possible for the segment-internal element and the connective to each signal part of the relation that should be inferred.

5.3.1 Division of labor between connectives and segment-internal elements

An obvious example of a *division of labor*-type of interaction between parts of the discourse segments and the connective can be found in relations with what Hoey (1983) labels ‘lexical signaling.’ Hoey (1983:44-53) argues that the relation in (8) can be, somewhat wordily, paraphrased as (9).

- (8) I beat off the attack **while** I was on sentry duty, by opening fire **when** I saw the enemy approaching.
- (9) The cause of my opening fire was that I saw the enemy approaching. The circumstances of my seeing the enemy approaching was that I was on sentry duty.

The coherence relations signaled by connectives in (8) are expressed by lexical items in (9). This eliminates the need for connectives, and vice versa; the presence of connectives in (8) renders the use of the lexical signals from (9) obsolete.

The fragment in (9) is a very contrived example. Actual instances of this type of signaling can be found in the Penn Discourse Treebank (PDTB; Prasad et al. 2008). In the absence of a connective, the PDTB instructs annotators to insert a connective that best expresses the inferred relation. However, in situations in which inserting a connective may lead to “redundancy in the expression of the relation ... because the relation is alternatively lexicalized by some ‘non-connective expression’,” annotators do not have to supply a connective (PDTB Research Group 2007:22). Instead, they are instructed to categorize the marking of the coherence relation as *AltLex* (Alternative Lexicalizations). An example of a RESULT relation marked by an alternative lexicalization can be found in (10).

- (10) [Ms. Bartlett’s previous work, which earned her an international reputation in the non- horticultural art world, often took gardens as its nominal subject.]_{S1} [Mayhap this metaphorical connection made the BPC Fine Arts Committee think she had a literal green thumb.]_{S2} (PDTB Research Group 2007:22-23)⁵

⁵ We copied the original PDTB segmentation.

In (10), the underlined alternative lexicalization indicates, with the verb *made*, that S_2 provides the result of S_1 , which is similar to what the function of, for instance, *so* would be. Adding *so* would be superfluous or result in a fragment of which the interpretation does not correspond to the current construction, for instance a conclusion. Some of PDTB's *Alternative Lexicalizations* could be considered more elaborate instances of connectives or cue phrases, for example *the most likely reason for this disparity is that* (PDTB Research Group 2007:23) being a more specific version of *the reason is that*. The crucial differences, however, are that *AltLexes* often include an anaphoric reference to the other segment (PDTB Research Group 2007) and that they often, as in (10), seem to contain propositional content.⁶ As such, many *AltLexes* are more appropriately considered as part of the segments than as the marking of the coherence relation only. *Alternative Lexicalizations* can thus be considered as segment-internal elements that convey information about the coherence relation that should be established, although each instance tends to be novel and can signal a very broad range of relation types.

A more specific segment-internal cue that has been linked to a particular type of coherence relations is *negation*. Using corpus data from the PDTB, Webber (2013) shows that the presence of negation markers in the first segment is a cue for CHOSEN ALTERNATIVE relations. In CHOSEN ALTERNATIVE relations, which are often signaled by *instead*, the two segments present alternatives, of which one is taken (PDTB Research Group 2007:36), as in (11).

- (11) [You cannot buy tickets online.]_{S1} **Instead**, [you have to make sure you get to the venue on time.]_{S2}

Webber (2013) argues that a negation marker suffices to indicate that one of the alternatives has been excluded. This does not only explain why negation is a common feature of CHOSEN ALTERNATIVE relations in general, but also why it is even more frequent in CHOSEN ALTERNATIVE relations that are not marked by a connective (Webber 2013). In addition to negation markers, Webber (2013) points toward other elements that create negative assertions (a subset of downward entailing constructions and modal markers that indicate that an event does not hold) as cues heavily associated with CHOSEN ALTERNATIVE relations. These cues are also more frequent in implicit than in explicit CHOSEN ALTERNATIVE relations. Since a negation element in itself signals that only one of the two segments of a relation holds, which is very similar to the function of *instead*, and since negation elements occur much more frequently in implicit CHOSEN ALTERNATIVE relations, it seems that in this combination of segment-internal cue and connective, the task of signaling the relation can be taken

⁶ Connectives are generally considered to have procedural meaning, rather than propositional meaning (e.g., Blakemore 1987, Wilson & Sperber 1993).

up by either one of these linguistic elements. As such, negation+*instead* can be considered an example of a *division of labor*-type of interaction.

Another example of an interaction between segment-internal features and connectives that appears to fall into the division of labor category is certain combinations of verb tenses and connectives specifying ORDER (e.g., Kehler 1994, Lascarides & Asher 1993).

- (12) [Hank was crying non-stop]_{S1} **after** [his girlfriend broke up with him.]_{S2}
 (13) [Hank was crying non-stop.]_{S1} (**after**) [His girlfriend had broken up with him.]_{S2}

In both (12) and (13), S₂ takes place before S₁. In (12), this is signaled by the connective. In (13), the pluperfect in S₂ explicitly places the event before the event in S₁, expressed in the simple past.⁷ While in (13) the connective *after* may be added, it is certainly not necessary to arrive at the same interpretation. Removing the connective from (12), however, changes the interpretation of the relation; in (14) the crying is more plausibly interpreted as preceding and probably even causing the break-up.

- (14) [Hank was crying non-stop.]_{S1} [His girlfriend broke up with him.]_{S2}

The combination of verb tenses in (13) signals that the relation has a non-basic order and reduces or even eliminates the need for a connective indicating non-basic order. Conversely, if a connective specifies that S₂ took place before S₁, both segments can be expressed in the same tense. A similar example in which a combination of verb tenses can make a connective superfluous is two segments expressed in the French *passé composé*, which has been claimed to encode sequential, or basic, temporal order, in combination with the connective *puis*, which indicates the same thing. Sentences that have both segments in the *passé composé* and are marked by *puis* are dispreferred over sentences in which the main verb tenses and the connective do not both explicitly encode sequentiality (Grisot & Blochowiak *submitted*).

Other examples of division of labor type of interactions are a semantic opposition between elements in both segments and contrastive connectives, as illustrated in (2), repeated here as (15), and focus markers such as *only* in S₁ and contrastive connectives, as illustrated in (16). *Only* in S₁ indicates that the content of this segment, unlike its alternatives, holds true, which is highly similar to the function that *but* would have in (16) if *only* were left out (Carlson 2014, Umbach 2005). Note that in both these examples, a less specific connective like *and* would work fine as

⁷ In both cases, it is likely that a causal relation is inferred on top of the temporal relation.

well. Both (15) and (16) are negative additive relations. In case the relations are marked by *and*, the connective signals that the BASIC OPERATION of the relation is additive, while the semantic opposition or the focus marker would indicate that the POLARITY of the relation is negative.

- (15) [Jack is a great kid.]_{S1} (**but**) [His sister is horrible.]_{S2}
(16) [Kathlyn only liked her uncle]_{S1} (**but**) [She didn't like any of her other family members.]_{S2}

In division of labor types of interactions between segment-internal features and connectives, both types of signals are in themselves able to explicitly mark the same primitives of a coherence relation and presence of both types of signals is redundant. This makes it distinctly different from the other two types of interactions, as we will demonstrate in Sections 5.4 and 5.5. We discuss the implications that the differences between the three types of interactions have for processing in Section 5.6.

5.3.2 Division of labor in translation

In division of labor types of interactions between connectives and segment-internal elements, the presence of the segment-internal feature can make the connective, or part of the connective, redundant, and vice versa. When it comes to translation, it can therefore be expected that such segment-internal features can appear in the translation *instead of* the connective. Alternatively, a segment-internal element can be used to signal part of the primitive values signaled by the original ST connective. In (17), the original English relation uses *also* as a connective to link two arguments together. The German translation does not include a connective, but uses *auch* 'also' as an NP modifier on the subject, which is obvious from the word order in S₂.⁸ Despite the different constructions, the original and the translation yield very similar interpretations. This suggests that, at least in some relations, NP modifiers signaling additivity can function as an explicit signal for additive relations, and that this signal can replace the need for an additive connective.

- (17) EN [Strict time-limits would clearly be very useful here.]_{S1} **Also**, [the idea of a complaints register accessible on the Internet links the last debate with this one rather nicely and empowers those who have difficulties in this area.]_{S2} {ep-99-01-13}

⁸ With German being a V2 language, the difference between *auch* as an NP modifier on the subject and *auch* as a connective can be clearly distinguished by the position of the verb. If *auch* is used as a connective, the verb appears between *auch* and the NP; if *auch* is an NP modifier, as in (17), the verb appears after *auch* and the rest of the subject.

DE [Es liegt auf der Hand, daß strikte Zeitvorgaben hier sehr nützlich wären.]_{S1} [Auch die Idee eines über das Internet zugänglichen Beschwerderegisters stellt eine recht passende Verbindung zwischen dieser und der vorangegangenen Debatte her und gibt denjenigen ein Mittel an die Hand, die auf diesem Gebiet Schwierigkeiten haben.]_{S2}
 ‘It is obvious that strict time-limits would be very useful here. The idea of a complaints register accessible on the Internet as well ...’

Conversely, a connective in the TT can be used to replace a segment-internal element in the ST. In (18), for example, the Dutch connective *zodat* ‘so that’ is a connective signaling that a PURPOSE relation should be inferred between the segments. The English original, however, uses the more general causal connective *because*. The difference between the segments of the ST and the TT is that the English original includes *we want*, which does not show up in S₂ of the Dutch translation. Interestingly, the meaning of *we want* exactly expresses the intentionality of the causal relation that distinguishes PURPOSE relations from other types of causal relations (e.g., Reese, Hunter, Asher, Denis, & Baldridge 2007). This example thus demonstrates that the combination of *because* and an expression of intentionality inside the antecedent can together signal a PURPOSE relation: an example of division of labor between a segment-internal element and a connective.

- (18) EN [Help us to be more precise with the road-map]_{S1} **because** [we want to follow it.]_{S2} {ep-00-04-11}
 NL [Help ons de routebeschrijving nauwkeuriger te maken,]_{S1} **zodat** [we hem kunnen volgen.]_{S2}
 ‘Help us make the road-map more precise, **so that** we can follow it.’

A highly frequent example of division of labor in the corpus can be found in English relations marked by *unless* (negative conditional relations). While all target languages possess a grammaticalized equivalent connective or cue phrase (DE *es sei denn*, ES *a menos que*, *a no ser que*, FR *à moins que*, NL *tenzij*), these are often not used to translate *unless*. Frequently, the translations make use of a connective equivalent to *if* (DE *wenn*, ES *si*, FR *si*, NL *als*). In these translations, a negation element is introduced in one of the segments (usually the antecedent), as in (19). If, however, a connective equivalent to *unless* is used, no negation elements are introduced inside the segments, as is illustrated by (20).

- (19) EN **Unless** [we take that way,]_{S1} [the only alternative will be more misery, more destruction and more death.]_{S2} {ep-02-04-09}
 FR **Si** [nous n'empruntons pas cette voie,]_{S1} [nous assisterons à plus de misère, plus de violence et plus de morts.]_{S2}
'If we don't take this path, we will witness more misery, more destruction and more death.'
- (20) EN [There is no reason for their disappearance]_{S1} **unless** [we condemn them to such a fate.]_{S2} {ep-01-12-17}
 FR [Il n'y a aucune raison qu'ils disparaissent]_{S1} **à moins que** [nous les condamnions à un tel sort.]_{S2}
'There is no reason they would disappear unless we condemn them to such a fate.'

Table 1 provides an overview of how often a negation element is inserted into a TT segment in translations with a connective equivalent to *unless* versus in translations with a connective equivalent to *if*. Negation is only added to a TT segment when the negative value for POLARITY is not expressed by the TT connective, as is the case with *if* ($\chi^2(1)=638.99$, $p<.001$).⁹ Since this trade-off pattern was the same in all four language pairs, we grouped all data together.

Table 1
Translations of unless into Dutch, German, French, and Spanish

	'Unless'	'If'
+ negation in TT	0	304
- negation in TT	339	0

A similar division of labor pattern can be found in the translation of English relations marked by *if* into, for instance, German. Subjunctive mood (*Konjunktiv II*) can be used to express conditionality, as is illustrated by (21).

- (21) EN **If** [there had been a check with cards,]_{S1} [maybe we could have avoided the problem of having a head count.]_{S2} {ep-00-04-13}
 DE [**Wäre** eine Überprüfung mit Karten vorgenommen worden,]_{S1} [hätten wir vielleicht das Problem des Abzählens der einzelnen Mitglieder umgehen können.]_{S2}
'Had a check with cards been made, we would perhaps have been able to avoid the problem of having a head count.'

⁹ All statistical tests were performed in R (R Core Team 2016, version 3.2.4).

Table 2 shows that if the German translation uses a subjunctive, it does not include a connective, and that when the translation uses a positive conditional connective equivalent to *if* (*wenn*), it does not introduce a subjunctive mood to the antecedent ($p < .001$, Fisher's exact test).

Table 2
Translations of if into German

	<i>Wenn</i> 'If'	No connective
+ subjunctive in TT	0	31
- subjunctive in TT	187	0

The combinations of segment-internal features and connectives discussed in this subsection are all clear examples of division of labor type of interactions. In translation, a feature can be used to replace the ST connective or part of it, and vice versa. This translation pattern is distinct from what can be seen in agreement and general collocation types of interactions, as will be illustrated by the data presented in the next two sections.

5.4 Agreement

Even if the meaning signaled by a segment-internal feature overlaps with the meaning signaled by a connective, this does not necessarily mean that the presence of one can make the presence of the other redundant, as is the case with division of labor type of interactions. In these combinations of segment-internal elements and connectives, the presence of the feature does not affect whether or not a connective is used to signal the same meaning. In these cases, there seems to be *agreement* between the feature and the connective.

5.4.1 Agreement between connectives and segment-internal elements

One example of an agreement type of interaction can be found in non-volitional causal relations in Dutch. Non-volitional causal relations feature a cause that does not involve a volitional agent, and a result, as in (22) (e.g., Mann & Thompson 1988, Stukker, Sanders, & Verhagen 2008).

- (22) **Because** [the airport's main runway was covered in snow,]_{S1} [all flights were delayed.]_{S2}
- (23) **Doordat** [er te veel sneeuw op de grootste vertrekbaan van het vliegveld lag,]_{S1} [hadden alle vluchten vertraging.]_{S2}

(23) is the Dutch equivalent of (22). In this example, the relation is marked by *doordat* ‘because of the fact that’ which is a positive causal connective specified for non-volitionality (Stukker et al. 2008). If the general positive causal connective *omdat* had been used, or even no connective, the relation would still have been interpreted as a non-volitional causal relation, since the presence of the non-volitional event in the antecedent is sufficient to signal the non-volitionality. By definition, all non-volitional causal relations contain a non-volitional event, and yet Dutch has a specific causal connective that appears in these contexts. This type of interaction is thus crucially different from the type of interaction between the segment-internal features and connectives found in the previous section. While both the segment-internal element and the connective signal a similar feature, they tend to appear together. In these cases, there does not seem to be a division of labor between the different signals, but rather agreement.

A similar type of agreement can be found in positive subjective causal relations. Such relations contain a conclusion, claim, or judgment by the speaker, with an accompanying motivation. Many consequents of subjective causal relations feature subjective words that convey the speaker’s attitude (e.g., *stupid*) or otherwise indicate that the content of the segment stems from the speaker’s mind (e.g., *obviously*, modal verbs), as in (24) (e.g., Canestrelli 2013, Wei 2018). However, this need not necessarily be the case, as is illustrated by (25).

- (24) [That guy must obviously be stupid,]_{S1} because [he tried to dry his socks in the microwave.]_{S2}
- (25) [Charlotte is dating someone,]_{S1} because [she has cancelled plans with us three times in the past two weeks.]_{S2}

In isolation, S₁ from (25) could just as well, if not more likely, be a fact. In S₂, however, a motivation is added, indicating that the relation involves the speaker’s reasoning, in which the first segment is treated as a claim. Several languages have been claimed to possess specific subjective causal connectives, which would be used to mark relations like the ones in (24) and (25). Examples of subjective causal connectives are Dutch *want*, German *denn* and *da*, French *car*, and Mandarin Chinese *jiran*, which can all be translated as ‘because’ (e.g., Li, Sanders, & Evers-Vermeul 2016, Pit 2003). Similar to non-volitional connectives relations marking non-volitional causal relations despite the presence of a strong non-volitional cue in the form of an event without an agent, subjective causal connectives are typically used to mark subjective causal relations even if the consequent contains subjective elements. In Dutch, for instance, the most prototypical connective in both (24) and (25) would be *want*; in German, it would be *denn* or *da*.

In Section 5.3.1, negation elements were discussed in relation to CHOSEN ALTERNATIVE relations, following Webber (2013). Asr and Demberg (2015) also look

at the presence of negation markers in coherence relations, but consider a wide range of relation types. Their findings are in line with Webber (2013) when it comes to CHOSEN ALTERNATIVE relations and even show that the presence of a negation marker is most strongly associated with this type of relation. In addition, they find that negation markers also often appear in other types of negative relations, for example in EXPECTATION, CONTRAST, COMPARISON, and CONTRA-EXPECTATION relations (all PDTB 2.0 labels). For most of these relations – all but EXPECTATION – however, it is not the case that negation markers appear more often in the implicit than in the explicit realizations of the relation. Since EXPECTATION, CONTRAST, COMPARISON, and CONTRA-EXPECTATION are all negative relations, the meaning signaled by the segment-internal negation element seems to overlap with the connectives prototypically used to mark these relations (*but, however, although, even though*); both the feature and the connective indicate a negative value for POLARITY. Unlike for CHOSEN ALTERNATIVE relations, the presence of a negation appears, by itself, not to be sufficient to signal these types of relations.¹⁰ Consider for instance (26) and (27), negative causal relations that when using PDTB 2.0 would be classified as EXPECTATION relations.

- (26) [Gary has never finished a knitting project,]_{S1} **even though** [he loves to knit]_{S2}
 (27) [Gary recently threw out all his knitting supplies,]_{S1} **even though** [he loves to knit]_{S2}

In both (26) and (27), S₂ denies an expectation that S₁ sets up, the underlying assumption being that people who never finish knitting projects or get rid of their knitting supplies do not love knitting. (27) demonstrates that a negation element is not required for negative causal relations, while (26) shows that it is also not sufficient to signal a negative causal relation. It does not even suffice to signal just the POLARITY of the relation, since substituting *even though* with *because* would result in a very different, somewhat incomprehensible fragment.

The interaction between subjective causal and non-volitional causal connectives and their corresponding segment-internal cues, as well as the interaction between negation elements and connectives marking EXPECTATION, CONTRAST and CONTRA-EXPECTATION relations, is thus crucially different from the interactions discussed in Section 5.4.2. Unlike in division of labor type of interactions, the presence of a segment-internal cue in agreement type of interactions does not seem to

¹⁰ It should be noted that COMPARISON is the most general type of negative relation in the PDTB 2.0 inventory and that CONTRAST includes many more types of relations than, for instance, the ones in (15) or (16). Many relations in the CONTRAST class are, for example, JUXTAPOSITIONS, which do not feature direct opposites, like *nice* and *horrible*, but rather non-identical alternatives, like *\$5* and *\$10*.

make the presence of a connective redundant, and vice versa, even though the two signals overlap in terms of the primitive values they explicitly encode. Although this type of interaction appears to go against what would be predicted on the basis of the linguistic theories discussed in Section 5.1.1, *agreement* is a very common linguistic phenomenon (e.g., number, gender, or case agreement), and it has often been demonstrated that language tends to be less economical, and thus more redundant, than minimally required (e.g., Bazzanella 2011, Horn 1993). Both observations also seem relevant to the marking of coherence relations; often, a connective agrees with the type of relation it occurs in (i.e., causal connectives usually appear in causal relations, temporal connectives most prototypically occur in temporal relations, etc.), and more coherence relations are explicitly marked than would be absolutely necessary, as is for instance illustrated by the fact that many of the examples in this chapter are perfectly acceptable both with and without connective, e.g., (22), (24), and (25). In Section 5.5, we will discuss a final type of interaction, *general collocation*, in which there is no overlap between the meaning encoded by the connective and the segment-internal cue.

5.4.2 Agreement in translation

As was argued in the previous section, an example of an agreement type of interaction can be found in subjective causal relations in Dutch. In general, the subjectivity profiles of Dutch causal connectives are very well studied. Dutch differs in its use of causal connectives from English. *Because*, the most frequent English causal connective used in non-basic causal relations, is widely used in both subjective and objective causal relations (Ford 1993, Sweetser 1990). Dutch, on the other hand, tends to use specific causal connectives depending on the subjectivity of the relation. *Want* is the most frequent subjective causal connective; objective relations are most frequently marked by *omdat* (overall, *omdat* is the most generic backward causal connective) (e.g., Sanders & Spooren 2015). When translating relations marked by *because* into Dutch, a choice has to be made between using a subjective or an objective causal connective. Although translators are probably not consciously aware of the difference in subjectivity between the Dutch causal connectives, it can be expected that subjective causal connectives will most often be used to translate subjective *because* relations and that objective *because* relations will be most often translated using an objective causal connective. This was indeed the case in our corpus (see Table 3, $\chi^2(1)=38.85, p<.001$).¹¹

¹¹ Although Table 3 shows that objective causal connectives tend to be used to translate *because* in objective relations and subjective causal connectives tend to translate *because* in subjective relations, the distribution of subjective and objective connectives over subjective and objective relations is probably not identical to patterns found in non-translated data (see also Cartoni, Zufferey, Meyer, & Popescu-Belis 2011, Degand 2004). Translations of *because* from English into Dutch are expected to be biased toward *omdat* (= objective), because it is syntactically equivalent to *because* (both subordinating conjunctions; *want* is a coordinating conjunction) and because it is the most frequent and most general Dutch connective.

Subjective causal relations tend to contain subjective words in their antecedent. We annotated the antecedents of all *because* relations in our corpus to determine the presence of subjective cues ($\kappa=.78$, $AC1=.83$).¹² The subjective *because* relations indeed often contained subjective words; the objective *because* relations usually did not contain subjective cues (see Table 3, $\chi^2(1)=202.38$, $p<.001$).¹³

Being an agreement-type of interaction, the subjective cues often appear in addition to a subjective causal connective, even though both the cue and the connective signal that the consequent is subjective, which makes it highly probable that the relation is subjective. In the Dutch translations in (28) and (29), for instance, the relation is marked by *want*, a subjective connective, even though it is already obvious that the relation is constructed in the speaker's mind from, respectively, *het is erg jammer* 'it is a great pity' or the modal verb *moeten* 'must' in S_1 .

- (28) EN [It is a great pity indeed that Commissioner Barnier has been unable to be present here this morning,]_{S1} **because** [this is a matter within his brief which is causing great concern not only in Scotland and Wales but in other parts of the Union.]_{S2} {ep-00-03-17}
- NL [Het is erg jammer dat commissaris Barnier hier vanmorgen niet kon zijn,]_{S1} **want** [dit is een kwestie uit zijn bevoegdhedenpakket die niet alleen Schotland en Wales, maar ook andere regio's uit de Unie grote zorgen baart.]_{S2}
- (29) EN [We must take the matter up in the Staff Regulations,]_{S1} **because** [it is an important point.]_{S2}
- NL [We moeten het punt opnemen in het Statuut van de ambtenaren van de Europese Gemeenschappen,]_{S1} **want** [het is een belangrijk punt.]_{S2}

For translation, an agreement type of interaction implies that, unlike in division of labor types of interactions, the segment-internal cue and the connective will not be used as substitutes for each other. More specifically, we do not expect the presence of a cue in the antecedent of the TT relation to influence whether a connective will be used and we should expect to see less variation in cues added or removed from the

¹² At this point, we did not distinguish between different types of subjective cues, but this may be worth exploring in the future. Wei (2018), for instance, shows that the subjective Mandarin Chinese connective *kejian* 'so' collocates more often with subjective elements that can be classified as indicating 'attitudinal stance' such as *importantly* (Conrad & Biber 2000) than with subjective markers of 'epistemic stance,' such as modal verbs.

¹³ The category of subjective causal relations in the corpus includes relations with 1st person evaluators in the present tense. Since this is a category of relations that is not prototypically subjective (prototypically, subjective relations have an implicit *subject of consciousness*, see Chapter 3), we also ran all analyses without this group of relations. Since the results remained the same, we did not exclude them from the dataset used in the final analysis.

segments between ST and TT. Table 3 presents an overview of the Dutch translations of all subjective *because* relations in the corpus. As predicted, we found that cues and connectives are indeed not used as substitutes for each other, and that ST and TT relations do not differ in whether or not their segments contain subjective cues, as is also illustrated by the examples in (28) and (29); in other words, all *because* relations with one or multiple subjective elements in their consequents also contained subjective elements in the TT (also in the consequents).

Table 3

Dutch translations of objective and subjective because relations with and without subjective elements in the consequent

		Objective causal connective ¹⁴	Subjective causal connective ¹⁵	Temporal / additive / no connective
Objective causal relation	+ subj. element in ST+TT	21	13	3
	– subj. element in ST+TT	98	19	24
Subjective causal relation	+ subj. element in ST+TT	70	90	33
	– subj. element in ST+TT	6	6	2

The presence of a subjective cue cannot predict whether the TT uses a causal connective ($p=1.00$, Fisher's exact test); subjective relations with a subjective element in their consequent receive a causal connective in the TL as often as subjective relations without a subjective cue. Finally, subjective relations with a subjective cue were translated using a subjective connective as often as subjective relations without a subjective cue ($\chi^2(1)=0.01$, $p=.90$). If the interaction between subjective elements and subjective causal relations had been division of labor, relations with a subjective cue would not be translated using a subjective causal connective.

Subjective causal relations thus show a translation pattern very different from the negative and positive conditional relations discussed in Section 5.3.2. These differences can be explained on the basis of the different types of interactions between segment-internal cues and connectives found in the respective relations. Both are in turn distinct from the final type of interaction we defined, general collocation, as will be shown in the next section.

¹⁴ *Omdat* 'because', *doordat* 'because (of the fact that)', *daar* 'for', *door het feit dat* 'because of the fact that'.

¹⁵ *Want* 'because/since', *immers* 'after all', *aangezien* 'considering', *namelijk* 'namely', *gezien het feit dat* 'considering the fact that'.

5.5 General collocation

There are many segment-internal features that have been associated with specific types of coherence relations in which the feature and the connective do not signal the same primitive values. Since connectives most prototypically associated with a specific type of coherence relation tend to signal primitive values inferred in those types of relations (e.g., positive causal connectives tend to mark positive causal relations), in this type of interaction the segment-internal elements do not seem to explicitly signal the relation at all. Instead, they seem to function as a cue mainly because they often co-occur with a specific type of relation. In other words, the segment-internal elements and the connectives/relations seem to be *general collocations*.

5.5.1 General collocation between connectives and segment-internal elements

Asr and Demberg (2015) find that negation elements are associated with, among relation types discussed in Section 5.4.1, REASON and RESULT relations. Here, the meaning of the segment-internal cue does not overlap with the meaning of the connectives that prototypically mark these types of relations (*because, so*), since REASON and RESULT relations both have a positive value for POLARITY. Rather than division of labor or agreement, the type of interaction between negation elements and the connectives associated with these relations can be more appropriately described as a *general collocation*; the segment-internal cue and the connective often appear together in a relation without being semantically related.

A general collocation type of interaction can also be found between causal relations and verb tense. Pit (2003) reports that subjective causal relations often appear in present or future tense, while objective causal relations more often connect segments in past tense. As such, past tense often co-occurs with objective causal connectives, such as German *weil* or French *parce que*; subjective causal connectives, such as Dutch *want* or German *denn*, co-occur more often with present or future tense.

Other examples of general collocations can be found in implicit causality (IC) verbs and verbs of transfer. Both types of verbs have been studied extensively when it comes to their effect on coreference patterns, with NP1 IC verbs (e.g., *apologize, disappoint*) preferring its subject to be the subject of the next clause and NP2 IC verbs (e.g., *admire, fire*) preferring its direct object to be the subject of the next clause (e.g., Au 1986, Garvey & Caramazza 1974, Koorneef & van Berkum 2006, Stewart, Pickering, & Sanford 2000); transfer verbs (e.g., *hand, give*) tend to prefer continuations about their goals (usually the indirect object), rather than their sources (Rohde, Kehler, & Elman 2006, Stevenson, Crawley, & Kleinman 1994). It has, however, been shown that the coreference patterns associated with both IC verbs and verbs of transfer are actually mediated by the type of coherence relation in which they occur. The coreference biases of IC verbs are contingent upon causal relations, either

basic or non-basic order positive objective causal relations, i.e., RESULT/CONSEQUENCE or REASON/EXPLANATION (e.g., Koornneef & Sanders 2013, Solstad & Bott 2013); the bias toward continuations about the goal of transfer verbs is mostly found in OCCASION relations, which are a specific type of basic order temporal relations (Rohde et al. 2006) or in RESULT relations (Stevenson et al. 1994).¹⁶ These types of relations are also the types of relations in which segments containing IC verbs or, respectively, verbs of transfer frequently occur. Implicit causality verbs thus function as a cue for positive objective causal relations, and transfer verbs for basic order temporal relations. Neither type of verb, however, explicitly encodes information about the type of relation that should be inferred between the segment in which it occurs and the upcoming segment, nor do they influence which connective is used to mark the relation. It thus seems likely that both verb types, as well as other segment-internal features in general collocation types of interactions, function as a signal for a specific type of coherence relation by virtue of frequent co-occurrence, i.e., collocation; language users are used to seeing the two together and, as such, encountering a segment-internal feature can help them predict or infer the upcoming coherence relation.

5.5.2 General collocation in translation

In general collocation types of interactions, there is no overlap in the primitive values signaled by the connective and the meaning expressed by the segment-internal element. As such, they cannot replace each other, and the presence of the one is not expected to make the presence of the other redundant. The translation patterns observed for coherence relations with division of labor and agreement types of interactions between their connectives and segment-internal elements are therefore not expected to be observed in general collocation types of interactions.

One of the general collocation types of interactions identified in Section 5.5.1 holds between negation elements and causal connectives. In translation, there should be very little variation in the presence of negation in the segments of the relation between the ST and the TT; if there is a negation in the ST, it is expected to also appear in the TT, and, conversely, if there is no negation in the segments of the ST relation, negation is not expected to appear in the translation. In addition, the presence of negation in the TT is not expected to result in an absence of the connective, or in the use of a connective that is less explicit than a causal connective (additive or temporal connective). (30) is an example of an English *because* relation with a negation element, along with its translation into all four languages in the corpus. All four translations include the negation and use a causal connective.

¹⁶ OCCASION: Infer a change of state for a system of entities in S_2 , establishing the initial state for this system from the final state of S_1 (see also Hobbs 1990).

- (30) EN [We were not very happy with the early versions from the rapporteur]_{S1}
because [we felt the targets were there to almost micro-manage the
market in renewable energy sources.]_{S2} {ep-00-03-29}
- DE [Mit den ersten Versionen des Berichterstatters waren wir nicht so
zufrieden,]_{S1} **da** [wir spürten, die Zielsetzungen waren dazu da, den
Markt erneuerbarer Energiequellen kleinzuhalten.]_{S2}
- ES [Las primeras versiones del ponente no nos gustaban demasiado]_{S1}
porque [teníamos la sensación de que el objetivo era gestionar al
milímetro el mercado de las fuentes de energía renovables.]_{S2}
- FR [Nous n'étions pas très satisfaits des premières versions du rapport]_{S1}
parce que [nous estimions que les objectifs n'existaient presque que
pour "microgérer" le marché des sources d'énergies renouvelables.]_{S2}
- NL [We waren niet erg ingenomen met de eerste versies van de
rapporteur,]_{S1} **omdat** [de doelstellingen erop gericht leken de markt
van hernieuwbare energiebronnen welhaast tot op het kleinste detail te
beheersen.]_{S2}

As can be seen from the overview of all translations of the *because* relations in the corpus in Table 4, translations never introduce or remove negation elements from the segments of the causal relation. There is also no difference between relations with or without negation when it comes to the use of a causal connective versus a less specific or informative connective option (either an underspecifying connective or no connective, $\chi^2(1)=0.02$, $p=.89$). Since this translation pattern did not differ between languages, we grouped all data together.

Table 4
Explicit vs. implicit translations of because relations with vs. without negation into Dutch, German, French, and Spanish

	'Because'	Temporal / Additive / No connective
+ negation in ST+TT	276	15
- negation in ST+TT	1171	69

Table 5 gives an overview of the connectives most frequently used to translate *because* into each language. Unlike for agreement types of interactions, in which the presence of a cue can function as a good predictor for the TT connective, as was shown in Section 5.4.2, causal relations with a negation element are not more associated with one connective over other candidate connectives, in any language ($p>.05$ for all contrasts). In other words, while there are causal connectives that tend to mark causal relations featuring subjective elements, there seem to be no causal connectives in

Dutch, German, French, or Spanish that are preferably used in causal relations featuring negation elements.

Table 5
Most frequent connectives ($n > 20$) used to translate because relations with vs. without negation into Dutch, German, French, and Spanish

	Dutch	n	German	n	French	n	Spanish	n
+ negation	<i>Omdat</i>	35	<i>Weil</i>	35	<i>Car</i>	28	<i>Porque</i>	59
in SL+TL	<i>Want</i>	18	<i>Denn</i>	21	<i>Parce que</i>	29	<i>Ya que</i>	3
	<i>Aangezien</i>	2	<i>Da</i>	12				
– negation	<i>Omdat</i>	154	<i>Weil</i>	124	<i>Car</i>	147	<i>Porque</i>	255
in SL+TL	<i>Want</i>	55	<i>Denn</i>	95	<i>Parce que</i>	109	<i>Ya que</i>	18
	<i>Aangezien</i>	21	<i>Da</i>	49				

Relations involving general collocation types of interactions do not show a specific translation pattern, as was expected on the basis of its characteristics. The way they behave in translation is, however, distinct from relations featuring division of labor or agreement types of interactions. This suggests that there are indeed three different ways in which segment-internal features and connectives can interact.

5.6 Discussion and conclusion

In this chapter, we aimed to develop a systematic way of categorizing segment-internal elements as signals of coherence relations. On the basis of the different ways in which elements inside the segments interact with connectives in the marking of coherence relations, we proposed a three-way distinction between division of labor, agreement, and general collocation. In division of labor types of interactions, the connective and the other signal overlap in the meaning they encode, and the presence of one is likely to make (part of) the other redundant; in agreement types of interactions, the connective and the other signal overlap in the meaning they encode, but they are commonly used in addition to each other. In general collocation types of interactions, there is no overlap in the meaning signaled by the connective and the other signal.

The way in which connectives and segment-internal elements interact in the meaning of coherence relations does not only provide new insights into the question of how coherence relations are marked, but also comments on the questions of how and why elements other than connectives can function as signals of coherence relations. In addition, the categorization proposed in this chapter has been shown to make meaningful predictions about the way in which connectives are translated.

In division of labor and agreement types of interactions, it is fairly obvious why segment-internal elements function as cues for coherence relations, since they

signal (part of) the same meaning the connective signals or, in case the connective is absent, would signal. In general collocations, the segment-internal element does not explicitly encode which coherence relation should be constructed, but seems to function as a cue because it often co-occurs with a specific type of relation (note that the same mechanism may, but need not be at play in division of labor and agreement types of interactions). This frequent co-occurrence may, however, not be coincidental; plausibly, the driving force behind both this frequent co-occurrence and the fact that a segment-internal element can function as a signal for coherence relations is the expectation of a specific type of coherence relation that is raised upon encountering a specific segment-internal element. Solstad and Bott (2013), for instance, propose that IC verbs carry an empty “explanatory slot” that has to be filled by information from the discourse. If an explanation for the action expressed by the IC verb has not yet been provided, the explanatory information is expected to follow the IC verb. Such a mechanism would not only explain why IC verbs frequently feature in causal relations, i.e., why speakers often produce this combination, but also why IC verbs can function as signals for causal relations, i.e., why listeners expect and/or look to infer causal relations after an IC verb.

Similar mechanisms can be thought of for the other examples of general collocations discussed in Section 5.5.1. Sentences involving a negation, for example, often express that something did not happen, which is most relevant if the expectation was that it was going to happen (e.g., Jordan 1988). Why the event did not take place, or why the speaker knows it did not may thus warrant some explanation. Alternatively, an utterance featuring a negation expresses that the speaker does not intend to do something or does not like something. Since such messages are often not what the listener wants to hear, they will often be accompanied by an explanation (e.g., Clayman 2002, Pomerantz & Heritage 2013). Finally, PROBLEM-SOLUTION relations, a specific type of causal relations, involve a negatively evaluated situation, i.e., the problem (e.g., Hoey 1983, Sanders & Noordman 2000, Sanders, Spooren, & Noordman 1993). As such, the segment expressing the problem often contains negation elements (Jordan 1984). In sum, there may be a reason for why a segment-internal feature and a specific type of coherence relation frequently co-occur. This underlying mechanism then does not only explain why speakers often produce the two elements together, but also provides an additional explanation as to why – on top of expectations formed on the basis of plain frequency of co-occurrence – segment-internal elements can function as cues for a coherence relation in general collocation types of interactions.

When it comes to the classification of interactions between segment-internal elements and connectives or coherence relations, a potential additional distinction could thus be made between collocations in which there is an underlying cognitive explanation for the frequent co-occurrence of a segment-internal element and a

relation type, and collocations in which there is not. Such a distinction, however, would require a lot more research into this phenomenon. In addition, it remains to be determined whether such a distinction would improve the explanatory power of the classification.

Throughout the chapter, we discussed many segment-internal elements that have been associated with specific types of coherence relations. The ways in which those observations were made, however, vary from monolingual corpus-based methods (e.g., Asr & Demberg 2015, Pit 2003, Webber 2013), experimental work (implicit causality, verbs-of-transfer), and theoretical explorations (e.g., Carlson 2014, Kehler 1994, Umbach 2005). This yields the question of how the signaling of coherence relations by segment-internal elements is most effectively studied. Corpus-based methods may present the most ecologically valid way of studying signals, but, as already discussed in Section 5.1.2, categorization of signals may be largely trivial since there is no obvious one-to-one mapping of signals onto relations. In addition, signals may not be frequent enough to study extensively or locate at all. Corpus-based methods will therefore need to be supplemented with theoretical and experimental work, since these methods allow researchers to zoom in on a single aspect of language.

In addition, while theoretical explorations and monolingual corpus studies are powerful tools for formulating hypotheses about language use, experimental methods are more equipped to demonstrate the cognitive plausibility of language models. A clear processing prediction that can be made on the basis of this chapter is, for instance, that in the presence of a segment-internal cue, connectives will be less beneficial to readers in division of labor types of interactions (in which case they could even be disruptive) than in agreement or general collocation types of interactions. Using different, complementary approaches can result in a comprehensive and complete overview of a linguistic phenomenon. In this chapter, we have demonstrated the usefulness of yet another approach, the use of translation corpora, in studying how coherence relations are signaled. Parallel corpora present a valuable additional tool to research segment-internal signals of coherence relations, especially when it comes to division of labor and agreement types of interactions; it is in these two types of interactions that translation can help make distinctions that in monolingual corpus data would largely be left to the interpretation of the researcher.

In general, relations that contain a segment-internal signal are hypothesized to less often contain a connective than relations that do not contain a segment-internal signal. However, as the current chapter argued, not all non-connective signals for coherence relations are created equal. The presence of a segment-internal signal is a much stronger predictor for the absence of the connective in division of labor types of interactions than in agreement or general collocation types of interactions. While the presence of a segment-internal signal may still increase the likelihood of more relations without a connective in agreement and general collocation types of interactions, the mechanism appears much weaker here and more susceptible to other

factors influencing the explicit versus implicit marking of coherence relations. Such factors for instance include the segment-internal signal being more strongly associated with another type of coherence relation (e.g., negation is a stronger cue for CHOSEN ALTERNATIVE relations than for causal relations, so a causal relation containing a negation may require a connective to block the expectation of a CHOSEN ALTERNATIVE relation), language users' default expectations about upcoming coherence relations (see Chapter 4), or the relations' position in the hierarchical discourse structure (see Chapter 4). Further examination of how different factors that appear to influence the marking of coherence relations by means of a connective work together seems imperative to fully understanding when language users use connectives to mark coherence relations in a discourse.

6 Restrictive relative clauses as discourse segments

Coherence relations are often assumed to hold between clauses, but restrictive relative clauses (RCs) are often not granted discourse segment status because they are syntactically and conceptually integrated in their matrix clauses. This chapter investigates whether coherence relations can be inferred between restrictive RCs and their matrix clauses. Four experiments provide converging evidence that suggests that restrictive RCs can indeed play a role at the discourse level and should not categorically be excluded from receiving discourse segment status in discourse segmentation and annotation practices. At the same time, the studies provide new insights into implicit causality verb biases, specifically about next-mention biases in concessive coherence relations, and discourse-level expectations about discourse structure, upcoming pronouns, and upcoming coherence relations.

The experiments in this chapter were conducted in collaboration with Hannah Rohde (University of Edinburgh)

6.1 Introduction

Inferring coherence relations is a prerequisite for understanding a discourse. A coherence relation can be defined as “an aspect of meaning of two or more discourse segments that cannot be described in terms of the meaning of the segments in isolation;” the meaning of a coherence relation is “more than the sum of its parts” (Sanders, Spooren, & Noordman 1992:2). A lot of attention has been paid to the types of coherence relations language users infer between discourse segments (e.g., Asher & Lascarides 2005, Carlson & Marcu 2001, Hobbs 1990, Kehler 2002, PDTB Research Group 2007, Reese, Hunter, Asher, Denis, & Baldridge 2007, Sanders et al. 1992, Wolf & Gibson 2005). Much less research has investigated between which parts of a discourse language users infer coherence relations (notable exceptions are Matthiessen & Thompson 1988, Polanyi 1988, Schilperoord & Verhagen 1998, and Verhagen 2001). Many approaches to discourse annotation have taken the grammatical clause as the basis for identifying discourse segments (e.g., Evers-Vermeul 2005, Mann & Thompson 1988, Sanders & van Wijk 1996, Wolf & Gibson 2005), but there tend to be exceptions to this rule; not all clauses can be discourse segments. One type of clause that is commonly excluded from receiving discourse segment status is the restrictive relative clause (RC) (e.g., Mann & Thompson 1988, Reese et al. 2007, Sanders & van Wijk 1996, Verhagen 2001).

Schilperoord and Verhagen (1998) use the notion of *conceptual dependency* to explain why restrictive RCs, as well as some other types of embedded clauses, such as clausal complements or restrictive adverbial clauses, are often excluded from being independent discourse segments, see (1).

- (1) If a constituent of clause A is conceptually dependent on a clause B, B is an integral part of the conceptualization of A, and therefore not available as a separate discourse segment (cannot enter into a discourse coherence relation with A, or any other part of the discourse).
(Schilperoord & Verhagen 1998:150)

Restrictive RCs syntactically link to a noun phrase (NP). Unlike non-restrictive RCs, they provide crucial information about the NP they modify, without which the conceptualization of the NP is incomplete. Clauses that contain a restrictive RC are therefore conceptually dependent on the RC and, as such, the RC and the matrix clause are assumed to form an integrated whole instead of independent discourse segments. While non-restrictive RCs, as in (2), are traditionally considered to be discourse segments, restrictive RCs, as in (3), thus seem to be excluded as discourse segments because they are both syntactically and conceptually integrated into their matrix clause.

- (2) Marcus, who is brilliant, now works at NASA.
- (3) A guy I knew in high school now works at NASA.

In some restrictive RC constructions, however, it is very plausible that a coherence relation is inferred between the restrictive RC and its matrix clause. Examples of such constructions are (4-6), as is also illustrated by the paraphrases in (4'-6').

- (4) Man who attacked jogger in Seattle park sentenced to prison.
- (4') Man is sentenced to prison because he attacked a jogger in Seattle park.
- (5) "I Did. Not. Plagiarize. That. Paper," Laura Pottsdam says of the paper that was almost entirely plagiarized.
- (5') "I Did. Not. Plagiarize. That. Paper," Laura Pottsdam says of her paper, even though it was almost entirely plagiarized.
- (6) Anyone who cares about food should be eating in Texas.
- (6') If you care about food, you should be eating in Texas.

Another indication that it is indeed possible, at least sometimes, to infer a coherence relation between a restrictive RC and its matrix clause, comes from translation. The dataset used in the parallel corpus study that formed the basis for Chapters 4 and 5 contained several examples of coherence relations that were translated by a restrictive RC construction, as in (7), and vice versa, as in (8). In both (7) and (8), the overall meaning of the translation is similar to the meaning of the original fragment. To indicate that in (5) the coherence relation in both languages holds within the complement of the attribution construction, the attribution has been underlined.

- (7) EN Recently we have seen headlines in Dutch and Irish newspapers about jet aircraft being chartered to fly workers from the west of Ireland to jobs in the Netherlands **because** the Netherlands cannot get workers to do this work. {ep-99-01-14}
- NL Onlangs meldden Nederlandse en Ierse kranten dat er vliegtuigen werden gecharterd om arbeiders uit het westen van Ierland naar Nederland te vervoeren voor banen **waar** geen Nederlandse werknemers voor kunnen worden gevonden.
"... for jobs **for which** no Dutch employees could be found."

- (8) EN However, those consular services are not available to Muslims from other EU Member States **who** would be there under the same terms and conditions operated by the Saudi authorities as UK Muslims.
{ep-01-01-17}
- DE Zu diesen Einrichtungen haben jedoch Muslime aus anderen EU-Mitgliedstaaten keinen Zugang, **obwohl** für diese dieselben Vorschriften der saudischen Behörden gelten, wie für die Muslime aus dem Vereinigten Königreich.
“... **although** they would be there under the same terms...”

A final source of evidence that suggests that restrictive RCs can enter into a coherence relation with their matrix clause can be found in Rohde, Levy, and Kehler (2011) and Kehler and Rohde (2015). In a continuation experiment, Rohde et al. (2011) asked participants to continue a relative clause that could be attached to two potential referents, see (9); in this example, *the children* and *the musician* compete for RC modification. In the experiment, the RCs were often used to supply an explanation for the main clause verb. Although the contents of the RCs were supplied by the participants (and restrictiveness can therefore not be guaranteed in all cases), the sample continuations Rohde et al. (2011:354-355) provide demonstrate that explanations can occur even when the RC is restrictive (for instance *Alan punished the accountant of the businessmen who was notorious for IRS fraud*).

- (9) John detests the children of the musician who ...

Rohde et al. (2011) also conducted a self-paced reading experiment in which participants read sentences such as the one in (9), completed with an RC that either modified the high NP (i.e., *the children*) or the low NP (i.e., *the musician*). The matrix clause verb was either a verb that raised the expectation of a causal relation featuring the direct object (i.e., an NP2-biased implicit causality verb, see Section 6.2) or a verb that did not. Since the high and low NP always differed in number, the finite verb of the relative clause disambiguated to whom the relative pronoun referred. Reading times revealed that high attachments were read slower than low attachments in the neutral verb condition, but faster in the causal expectation condition. This suggests that participants consider restrictive RCs a plausible location for causal information.

In the continuation experiment conducted by Kehler and Rohde (2015), participants supplied fewer explanations after a restrictive RC if a causal relation could be inferred between the RC and its matrix clause, as in (10), than in prompts where the RC merely provided additional information about its referent, as in (11).

- (10) The boss fired the employee who was embezzling money.

- (11) The boss fired the employee who was hired in 2002.

The experiments reported in Rohde et al. (2011) and Kehler & Rohde (2015) thus suggest that language users can infer causal relations between restrictive RCs and their matrix clauses.

The examples, translations, and experiments discussed above seem to provide evidence against the assumption that restrictive RCs cannot enter into a coherence relation with their matrix clauses. This chapter will further investigate this issue. It first explores (Experiment 1 and 2) whether a restrictive RC can enter into a relation with its matrix clause beyond the causal (explanation) relations tested previously, as is suggested by examples (2-6). Since conditional interpretations, such as the one in (6), seem most plausible in contexts that contain a quantifier, we focus mainly on exploring the availability of negative causal relations (also called denial of expectation relations); for ease of reference, these relations will be referred to as ‘concessive’ relations.¹

Experiment 1 and 2, like Experiment 1 in Rohde et al. (2011) and the experiment in Kehler and Rohde (2015), both use off-line measures to explore the possibility of coherence relations between restrictive RCs and their matrix clauses. Experiment 3 uses a self-paced reading paradigm to investigate whether restrictive RCs constructions show similar behavior in on-line processing as other types coherence relations between independent clauses. This experiment asks whether language users mainly make discourse-level inferences while reading restrictive RCs when they have to resolve syntactic ambiguity, as in Rohde et al. (2011), or actively engage with the RCs in an experimental paradigm because they have to provide a sensible continuation (as in Kehler and Rohde (2015), Experiment 1 in Rohde et al. (2011), and Experiment 1 and 2 of the current chapter), or if inferring coherence relations between restrictive RCs and their matrix clauses is a process that also occurs outside of these contexts. Finally, Experiment 4 uses an eye-tracking-while-reading paradigm to explore how restrictive RCs influence discourse-level expectations during on-line processing. Off-line experiments find that restrictive RCs can affect expectations about upcoming coherence relations (Kehler & Rohde 2015) and upcoming referents (Experiment 1 and 2). Eye-tracking allows us to more closely investigate how expectations are updated in real time.

While the distinction between restrictive and non-restrictive RCs is not always entirely clear-cut (e.g., Bache & Jakobsen 1980), we designed the RCs in all four experiments to be more characteristic of restrictive RCs than of non-restrictive RCs, using criteria listed in Bache and Jakobsen (1980) and Fabb (1990), among others: the RC is not separated from its matrix clause by means of a comma, the matrix clause

¹ The use of ‘concessive’ as a short-hand for negative causal relations should not be confused with the *concession* relations in Chapter 3, which are negative additive subjective relations.

and the RC can plausibly be uttered as a single intonation unit, the relative pronoun *who* can plausibly be substituted with *that*, and the RC cannot be removed from the sentence without losing essential information. In addition, the RC modifies the antecedent of the relative pronoun so that it refers to a unique referent (Experiment 1 and 2) or picks out a unique referent from a mentioned or invoked set of possible referents (Experiment 3 and 4).

All experiments reported in this chapter make use of implicit causality (IC) verbs, a well-studied linguistic phenomenon. Experiments 1, 2, and 4 make use of the finding that IC verbs generate expectations about the upcoming discourse and test whether the type of coherence relation between a restrictive RC and its matrix clause containing an IC verb has an impact on these expectations. A brief overview of the phenomenon of implicit causality will be given in Section 6.2, before Experiments 1-4 are reported in Sections 6.3-6.6.

Over the course of four experiments, this chapter thus tests the possibility of inferring a coherence relation other than a causal relation between a restrictive RCs and its matrix clause, whether the inference of coherence relations in restrictive RC constructions occurs even in contexts without ambiguity or specific task demands, and how quickly these discourse-level inferences occur. The experiments provide converging evidence that suggests that restrictive RCs can indeed have a function at the discourse level and should not categorically be excluded from receiving discourse segment status in discourse segmentation and annotation practices. At the same time, the studies provide new insights into IC verb biases, specifically about next-mention biases in concessive coherence relations, and discourse-level expectations about discourse structure, upcoming pronouns, and upcoming coherence relations.

6.2 Implicit causality

Implicit causality has been the topic of many psycholinguistic studies since it was originally described by Garvey and Caramazza in 1974 (Au 1986, Ferstl, Garnham, & Manouilidou 2011, Kehler, Kertz, Rohde, & Elman 2008, Koornneef & Sanders 2013, Koornneef & van Berkum 2006, Mak & Sanders 2013, McKoon, Greene, & Ratcliff 1993, among many others). IC verbs are transitive verbs that have a strong preference concerning the referent of the subsequent entity mention. NP1-biased IC verbs favor the continuing discourse to focus on their subject, as in (12); NP2-biased IC verbs favor the continuing discourse to focus on their object, as in (13).

- (12) **Tracy** annoyed Tom because **she** kept complaining.
- (13) Tracy fired **Tom** because **he** kept complaining.

The *next-mention bias* of IC verbs is specifically found in explanations of the event denoted by the IC verb (e.g., Hartshorne 2014, Kehler et al. 2008, Pickering & Majid 2007, Solstad & Bott 2013).

Another property of IC verbs crucial to the experiments reported in this chapter is that they raise the expectation of an upcoming explanation. In a continuation task, Kehler et al. (2008), show that IC verbs receive about 60% explanation continuations, while only 24% of continuations following non-IC verbs constitute an explanation.

There has been a lot of debate about what exactly gives rise to implicit causality biases, but most accounts seem to conclude that the biases are the product of (pragmatic) inferences (Hartshorne 2014, Kehler et al. 2008, Pickering & Majid 2007, among others). If these discourse-level inferences are influenced by restrictive RCs that can be related to their matrix clauses in a way that is more informative than merely providing information about one of the referents, and if this influence differs depending on the type of relation that can be inferred between the RC and its matrix clause, this would indicate that language users can indeed infer a coherence relation between restrictive RCs and their matrix clauses.

6.3 Continuation experiment 1

Experiment 1 tests whether restrictive RCs can influence next-mention expectations of the subsequent sentence. It aims to replicate the finding by Kehler and Rohde that restrictive RCs that also provide cause for the matrix clause event influence next-mention expectations, and to investigate whether restrictive RCs that also provide an unexpected cause (i.e., concessive RCs) can do the same. In this study, we presented participants with prompts for which they were asked to supply a natural continuation. Target prompts consisted of a matrix clause containing an NP2 IC verb, an object modified by an RC, and a connective. Prompts differed in the coherence relation that could be inferred between the RC and the matrix clause (causal, concessive, or neutral), and in the connective (*because* or *even though*), see (14). The full list of target prompts can be found in Appendix B.

- (14) We thanked the neighbor who . . .
- a. brought over a fruit basket because / even though . . .
 - b. dropped our newly inherited vase because / even though . . .
 - c. stopped by on Tuesday night because / even though . . .

In (14), each NP2 IC verb construction includes a restrictive RC modifying the object. If a causal relation is inferred between the restrictive RC and the main clause, as in (14a), the IC bias (i.e., an explanation featuring the NP2) should be fulfilled (Kehler & Rohde 2015). It can thus be expected that in the causal+*because* condition, the NP2 bias is reduced as compared to the neutral+*because* condition in (14c).

The prediction is a bit more complex for the concessive+*because* condition. A concessive relation between the restrictive RC and its matrix clause, as in (14b), indicates that something unexpected happens; thanking someone for ruining an heirloom is not a standard event. This discrepancy warrants an explanation. Compared to a neutral NP2 IC verb construction, as in (14c), there are multiple relevant candidates to focus an explanation on; the explanation may focus on the NP2 (e.g., *they offered to replace it*), but also on the NP1 (e.g., *we are too nice for our own good*), or on some other factor (e.g., *the vase may have been incredibly ugly*). Concessive RCs may thus reduce the NP2 bias for subsequent clauses, although not necessarily to the same extent as causal RCs.

The main reason why predictions about next-mention biases after a concessive RC are less straightforward than predictions for the causal condition is that much less is known about next-mention expectations after an IC verb in a concessive context than after IC verbs that feature in a causal relation. Several studies have explored IC biases after *but* (e.g., Ehrlich 1980, Koornneef & Sanders 2013), but we have not found any papers that specifically deal with IC verbs in concessive or denial-of-expectation relations. We added the *even though* conditions to investigate the effect that a concessive context has on next-mention biases. In the neutral+*even though* condition, there is expected to be no influence of the RC on the next-mention bias; this condition thus serves as a baseline condition to check the effect of a concessive relation on next-mention biases. We predict the NP2 bias in this condition to be reduced as compared to the neutral+*because* condition; *even though* signals that there is something unexpected going on, which, in the context of an NP2 biased IC verb boosts the relevance of the NP1, since they are doing something unusual. We expect the NP2 bias to be even further reduced for the causal+*even though* and concessive+*even though* conditions. With the causal RC (*We thanked the neighbor who brought over a fruit basket even though...*), the IC event in the matrix clause is explained reasonably by the content of the causal RC so we already can infer why the neighbor is being thanked but we don't know yet what could undermine that account of the events and either referent seems potentially relevant (... *even though we dislike her, even though she brought the gift for ulterior motives*). Finally, with the concessive RC, an unexpected state of affairs is presented (*We thanked the neighbor who dropped our vase...*); one possibility is that the connective *even though* will provide an opportunity to deal with this unexpected situation and contrast it with a different more normal scenario (...*even though we should have scolded her*) or it may provide an opportunity to clarify the nature of NP2's behavior that could more aptly link the behavior to NP1's thanking (...*even though she didn't mean to break it*). Across all of these predictions, the effect of *even though* appears to be to open up more possible avenues for subsequent coreference and to undermine the specific NP2 bias associated with the verb.

6.3.1 Participants

56 monolingual English speakers were recruited through Amazon Mechanical Turk (mean age 34.88, age range 23-66, 20 women). They participated in exchange for monetary compensation (\$6.00).

6.3.2 Materials

Participants were presented with 30 prompts consisting of a matrix clause containing an NP2 IC verb, an object modified by an RC, and a connective, see (14). The IC verbs used in all experiments reported in this chapter were taken from existing inventories of IC verbs (Commandeur 2010, Ferstl, et al. 2010, Koornneef & van Berkum 2006). The subject of all stimuli was a proper name or a first person pronoun; the direct object was a definite NP that specified (e.g., *guy*) or implied (e.g., a gardener is usually male) a different gender or person than the subject; the subject was male in 50% of the items with proper name subjects, female in the other 50%. The NP2-biased IC verb was always presented in the past tense. The full list of target items can be found in Appendix B.

The target prompts were distributed over six lists, with each item occurring only once per list, in one of the six conditions. Target prompts were interspersed with 16 fillers containing a connective, an embedded clause, or both, and 24 fillers from an unrelated experiment. The items from each list were presented to the participants in random order. Each participant saw every item only once, in one of the conditions.

6.3.3 Procedure

Continuations were collected via a web-based interface embedded in the Amazon Mechanical Turk environment. Each item was displayed on a separate page. Participants were instructed to write a natural continuation for the prompts in the supplied text box. Beforehand, the participants were informed that the experiment would not take longer than an hour; on average, participants took approximately 45 minutes to complete the experiment.

6.3.4 Annotation and data clean-up

One trained coder (author) annotated all continuations for the referent of the subject of the continuation: NP1, as in (15a), NP2, as in (15b), or Other, as in (15c).

- (15) Natalie distrusted the doctor who had messed up the procedure last time because
- a. **she** could have died.
 - b. **he** didn't seem to own up to his previous mistake.
 - c. **such a breach** of trust was hard to shake.

Being a relatively simple task, this type of annotation can generally be done very reliably, especially in contexts where the referents have different genders. We double-coded (author and another trained annotator) the continuations for Experiment 2, which replicate the next-mention results of Experiment 1. The annotation of Experiment 2 indeed shows a very high agreement for next-mention at 96%, $\kappa=.94$.

During annotation, we removed any unfinished continuations, as well as continuations that were completely nonsensical (4.11%). We only included continuations in which the connective attached to the main clause (87%) in our analysis (high attachment, see Section 6.4).

6.3.5 Analysis method

All experiments in this chapter were analyzed using linear mixed effects regression models (LMER; Baayen 2008, Baayen, Davidson, & Bates 2008) or, in case of categorical dependent variables, generalized linear mixed effects regression models (GLMM), using the lme4 package (Bates, Maechler, Bolker, & Walker 2015) in R (R Development Core Team 2016). For each model, we started with a maximal random effects structure, only simplifying the model in case of nonconvergence (Barr, Levy, Scheepers, & Tily 2013). We first reduced the random effects by taking out one or all correlations between random slopes and random intercepts. If the model would still not converge, we removed random intercepts or random slopes until we ended up with a converging model. See Barr et al. (2013) for a detailed account of this step-wise procedure.

The significance of fixed effects was determined by performing likelihood ratio tests to compare the fit of the model to that of a model with the same random effects structure that did not include the fixed effect. The categorical predictor variables in all analyses were deviation coded and centered. All pairwise comparisons were obtained using a subset of the data that only contained the relevant conditions with re-centered predictor variables.

6.3.6 Results

The proportions of NP1, NP2, and Other continuations per condition can be found in Figure 1. In our analysis, we modeled the binary outcome of NP2 versus non-NP2 continuations in a generalized mixed effects model. A likelihood ratio test found a significant interaction between RC type and connective ($\chi^2(1)=9.86, p<.01$; the model did not include random slopes for item and no random slope of connective for participant).

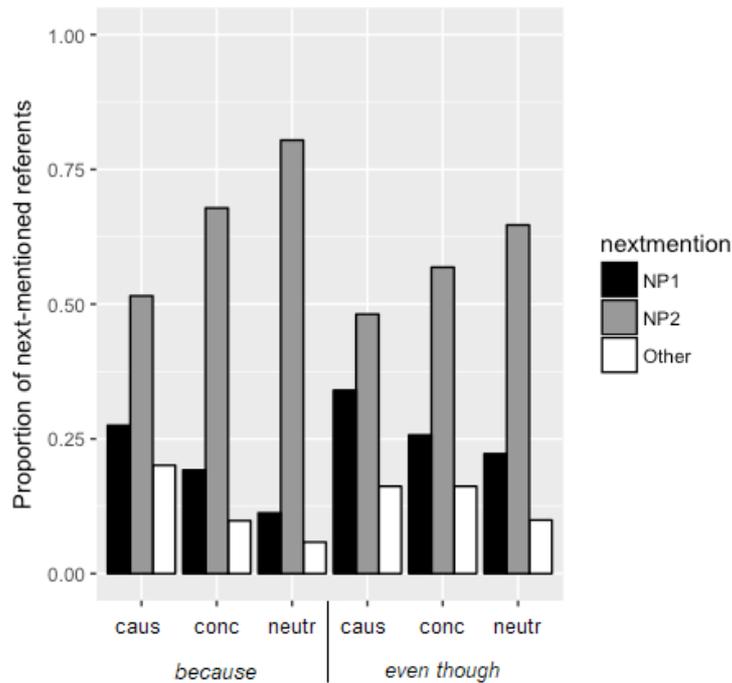


Figure 1. Proportion of next-mentioned referents per connective and RC type

Pairwise comparisons showed that there were fewer NP2 continuations after *even though* than after *because* in the neutral RC condition ($\beta=-1.33$ SE=0.38, $z=-3.49$, $p<.001$; the model did not include a random slope for item), the base-line condition where the RC should not influence the next-mention bias. There were also fewer NP2 continuations after *even though* than after *because* in the concessive RC condition ($\beta=-0.58$, SE=0.26, $z=-2.23$, $p<.05$; the model did not include random slopes for participant and item). There was no main effect of connective after causal RCs ($\beta=-0.31$, SE=0.36, $z=-0.85$, $p=.39$).

Analysis of the *because* subset revealed a main effect of RC type ($\chi^2(1)=17.18$, $p<.001$). Further comparisons between RC types within the *because* conditions revealed that there were fewer NP2 continuations after a causal RC than after a neutral RC ($\beta=-1.79$ SE=0.47, $z=-3.82$, $p<.001$) or after a concessive RC ($\beta=-0.80$ SE=0.33, $z=-2.43$, $p<.05$; the model did not include a random slope for participant). There were also fewer NP2 continuations after a concessive RC than after a neutral RC ($\beta=-0.77$ SE=0.37, $z=2.09$, $p<.05$; the model did not include a random slope for participant).

6.3.7 Discussion

The results from the continuation study show that the coherence relation between restrictive RCs and their matrix clauses influences IC biases. Causal RCs lead to the strongest reduction of the next-mention bias, but concessive RCs also reduce the proportion of references to the NP2. Similar to concessive RCs, a concessive connective reduces the next-mention bias of IC verbs.

There were not just fewer NP2 continuations after *even though* than after *because* in the neutral RC conditions, but also in the concessive RC conditions. This could be interpreted as an indication that the influence of the RC and the connective on the next-mention bias is cumulative; a concessive RC reduces the expectation of the NP2 being mentioned as the subject of the subsequent clause, after which the concessive connective further reduces this NP2 bias. The difference in the proportion of NP2 continuations after *even though* and *because* in the causal conditions, however, was not significant. This could potentially be due to the NP2 bias already being strongly reduced by the causal RC. For the causal+*because* items, approximately only half of all continuations had the NP2 as their subject.

In sum, the continuation study shows distinct next-mention patterns after causal, concessive, and neutral RCs. This suggests that language users indeed infer coherence relations between restrictive RCs and their matrix clauses.

6.4 Continuation experiment 2

Experiment 1 tested whether restrictive RCs can influence expectations about upcoming referents. When coding the continuations, we noticed that not all continuations attached to the main clause; some participants constructed a coherence relation between the continuation and the contents of the RC. The difference between these two constructions is illustrated by (16), where *because* attaches to the main clause, and (17), where the clause following *because* provides an explanation for the contents of the RC. As such, (16) and (17) have distinct discourse structures, respectively [MATRIX + RC] *because* [EXPLANATION] versus [MATRIX [RC] *because* [EXPLANATION]]. We will refer to attachments to the main clause as high attachments, and to attachments within the RC as low attachments.

- (16) **Mrs. Thompson loathed the gardener** who never took off his muddy shoes **because he made a mess all throughout her beautiful home.**
- (17) Mrs. Thompson loathed the gardener **who never took off his muddy shoes because he was embarrassed of his foot odor.**

Experiment 2 asks whether restrictive RCs can guide expectations about discourse structure. We hypothesize that if a causal relation is inferred between the RC and its matrix clause, there would no longer be an expectation for upcoming causal

information to explain the IC verb event. We would then expect any further causal cues to favor attachment to another part of the discourse, for instance the RC, compared to when the IC causal requirement has not yet been fulfilled. In other words, we expect fewer high attachments of *because* after a causal RC than after a neutral RC.

Since our discourse structure predictions apply specifically to explanation contexts, the prompts used in this experiment only include *because* as a connective. We do not necessarily expect concessive RCs to impact the discourse structure differently than the neutral RCs, but we kept the concessive RC condition to check if the next-mention results from Experiment 1 are replicable.

6.4.1 Participants

55 monolingual English speakers were recruited through Amazon Mechanical Turk (mean age 38.25, age range 22-67, 26 women). They participated in exchange for monetary compensation (\$5.50).

6.4.2 Materials

The target items in Experiment 2 were the same as the items in Experiment 1, with the exception of the connective manipulation; only the *because* versions were included in the second continuation experiment. 24 fillers were created to substitute the fillers from the unrelated experiment in Experiment 1. The new fillers all contained an embedded structure. Half of the items required or were biased toward high attachment, i.e., an attachment of the continuation to the matrix clause; the other half required or were biased toward low attachment, i.e., a continuation within the embedded clause, see (18). This manipulation was intended to avoid having an experimental bias toward high or low attachments. In Experiment 1, all of the fillers had been biased toward high attachment.

- (18) [high] Wade insisted that penguins did not really exist so ...
[low] The sales person guaranteed that we would get a full
refund if ...

The target prompts were distributed over three lists, with each item occurring only once per list, in one of the three conditions. Target prompts were interspersed with the attachment fillers and 16 additional fillers of various types. The items from each list were presented to the participants in random order. The experiment was conducted in the same way as Experiment 1, see Section 6.3.3.

6.4.3 Annotation and data clean-up

Two trained coders (author and an undergraduate Linguistics student) annotated all continuations for the referent of the subject of the continuation, using the categories

NP1, NP2, and Other (see also Section 6.3.4). The agreement between the coders was very high: 96%, $\kappa=.94$ (AC1=.94).

In addition, it was annotated for each continuation whether it attached to the matrix clause, as in (19a) or to the RC, as in (19b). We also included a label “both” for continuations that could plausibly be attached to the matrix clause as to the RC, or to both at the same time, as in (19c).²

- (19) Geoff ridiculed the stewardess who crashed the drink cart into the wall because ...
- a. he felt she was incompetent.
 - b. she was dizzy during the flight.
 - c. she was so clumsy.

Annotating the attachment of a continuation is a much more complex task than annotating co-reference in contexts with referents of different genders, since determining attachment relies more heavily on interpretation. Agreement between the two coders was satisfactory at 94% and $\kappa=.74$ (AC1=.96). Disagreements were resolved through discussion.

We removed any unfinished continuations, as well as continuations that were completely nonsensical (1.27%). We included only high attachments in our next-mention analysis (89%); low attachment continuations attach to the RC and, as such, the relation marked by *because* does not contain the IC verb in its relational segments. In attachments coded as ‘both’ the IC verb is included in one of the two relations marked by *because*, but since this discourse structure is distinctly different from the discourse structure found in high attachments, we excluded them from our next-mention analysis to keep our dataset as homogeneous as possible.

6.4.4 Results

When it comes to next-mention, we replicated the results from the *because* conditions from Experiment 1, see Figure 2. We used generalized linear mixed effects regression to model the binary outcome of NP2 versus not-NP2 continuations and found a main effect of condition ($\chi^2(1)=19.51, p<.001$). Pairwise comparisons revealed that there were fewer NP2 continuations in the causal condition than in the concessive condition ($\beta=-0.97, SE=0.30, z=-3.18, p<.01$) and the neutral condition ($\beta=-1.85, SE=0.34, z=-5.44, p<.001$). In addition, there were fewer NP2 continuations in the concessive condition than in the neutral condition ($\beta=-0.85, SE=0.35, z=-2.41, p<.05$).

² In cases where the content of the continuation appears to relate to both the RC and the matrix clause, the continuation seems to syntactically attach within the RC, but to conceptually reinforce the relation between the RC and the matrix clause.

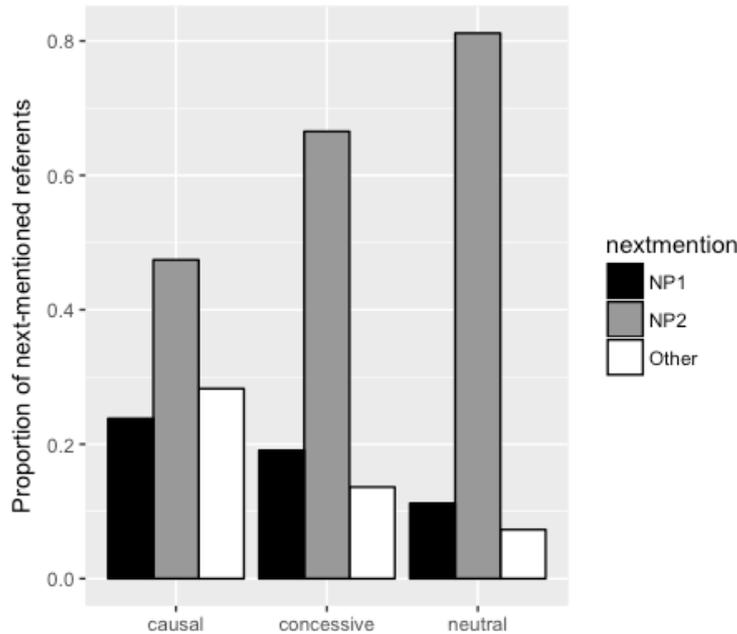


Figure 2. Proportion of next-mentioned referents per RC type

Figure 3 shows the proportion of high, low, and ‘both’ attachments. We analyzed the binary outcome of high versus not-high attachments using generalized linear mixed effects regression modeling. A likelihood ratio test confirmed a main effect of condition ($\chi^2(1)=30.15, p<.001$). Pairwise comparisons revealed that there were fewer high attachments in the causal condition than in the concessive condition ($\beta=-2.25, SE=0.61, z=-3.68, p<.001$) and the neutral condition ($\beta=-2.74, SE=0.40, z=-6.77, p<.001$; the model did not include random slopes for item and participant). There were also fewer high attachments in the concessive condition than in the neutral condition ($\beta=-1.34, SE=0.50, z=-2.67, p<.01$; the model did not include random slopes for item and participant).

We also modeled the binary outcome of low versus not-low attachments. A likelihood ratio test revealed a main effect of condition ($\chi^2(1)=658.40, p<.001$). Pairwise comparisons indicated that there were more low attachments in the causal condition than in the concessive ($\beta=9.73, SE=4.70, z=2.07, p<.05$; the model did not include random slopes for item) and neutral condition ($\beta=2.74, SE=0.79, z=3.46, p<.001$). There were also more low attachments in the concessive condition than in the neutral condition ($\beta=42.92, SE=8.89, z=4.83, p<.001$; the model did not include random slopes for item and participant).

Finally, we modeled the binary outcome of ‘both’ versus not-both attachments. A likelihood ratio test revealed a main effect of condition ($\chi^2(1)=45.70$, $p<.001$; the model contained no random slopes for item and participant). Pairwise comparisons indicated that there were more ‘both’ attachments in the causal condition than in the concessive ($\beta=2.51$, $SE=0.76$, $z=3.29$, $p<.01$) and neutral condition ($\beta=3.43$, $SE=1.08$, $z=3.18$, $p<.01$; the model contained no random slopes for participant). There was no difference in the proportion of ‘both’ attachments between the concessive and the neutral condition ($\chi^2(1)=1.72$, $p=.19$; the model did not include random slopes for participant).

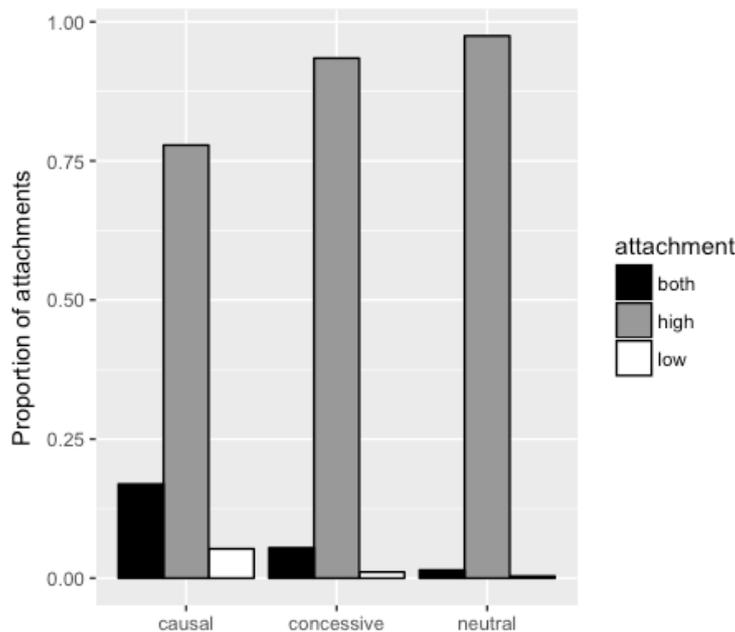


Figure 3. Proportion of attachments per RC type

6.4.5 Discussion

The next-mention results in Experiment 2 replicated the next-mention results from the *because* conditions in Experiment 1. The attachment results are in line with the prediction that there would be fewer high attachments in the causal condition than in the other two conditions. However, we expected this effect to be mainly driven by an increase in the number of low attachments in the causal condition. Even though the results do indicate that participants provided more low attachments in the causal condition than in the concessive or the neutral condition, there was a much bigger increase in the number of continuations that related to both the matrix clause and the

RC, as can be seen in Figure 3. In addition, we did not predict a difference in the proportion of high attachments between the concessive and neutral condition, but in the concessive condition we also found a decrease in the number of high attachments, as well as an increase in the proportion of low attachments.

Overall, the results from Experiment 2 are not entirely in line with our hypothesis that in a context in which the event of the matrix clause has received an explanation, readers are more inclined to expect a coherence relation within a restrictive RC. A more plausible explanation for our findings is that the inference of a coherence relation between a restrictive RC and its matrix clause increases the relevance of the RC at the discourse level; subsequent linguistic context is more likely to refer back to a restrictive RC if there was a coherence relation between the RC and its matrix clause. With distinct attachment patterns for each condition, however, the results do support our more global hypothesis that language users do infer coherence relations between restrictive RCs and their matrix clauses.

6.5 Experiment 3: Self-paced reading

Experiment 1 and 2 investigated the influence of restrictive RCs on expectations about the continuation of the discourse. The results of both studies corroborate the observation that readers can infer causal coherence relations between restrictive RCs and their matrix clauses, and established that other types of coherence relations are available as well. However, both experiments use off-line measures. As such, they provide no information about the time course with which the coherence relations are inferred; the results from Experiment 1 and 2 are not informative regarding whether or not a restrictive RC is a location where language users expect discourse-level information. This question is addressed in Experiment 3. By means of a self-paced reading task, we investigate whether the processing of coherence relations that hold between restrictive RCs and their matrix clauses mirrors the processing of coherence relations in more traditional constructions, e.g., between two independent clauses. The self-paced reading experiment in Rohde et al. (2011) already provides evidence that language users can expect restrictive RCs to convey causal information (see also Section 6.1). The current experiment investigates the availability of not just causal RCs, but also concessive RCs. In addition, the items from the current experiment do not involve the disambiguation of the relative pronoun; as such, this experiment examines whether the Rohde et al. (2011) findings were mainly due to participants using cues from the discourse to help disambiguate the reference of the relative pronoun, or whether restrictive RCs are generally places where language users expect discourse-level information.

A well-established finding in discourse processing is that causal information is processed faster than non-causal information, and that stronger causal links result in even faster processing times than weaker causal links (e.g., Haberlandt & Bingham 1978, Keenan, Baillet, & Brown 1984, Myers, Shinjo, & Duffy 1987, Sanders &

Noordman 2000, Wolfe, Magliano, & Larsen 2005). By contrast, several studies report slower reading times on relations with a negative value for POLARITY (i.e., relations involving some form of contrast, such as adversative, concessive or contrastive relations) than relations with a positive value for POLARITY (e.g., Clark 1974, Murray 1997, Wason & Johnson-Laird 1972). These findings also seem to hold true in the context of IC verbs; causal relations after an IC verb leads to faster reading times than additive or negative relations (Koorneef & Sanders 2013, Mak & Sanders 2013).

If coherence relations between restrictive RCs and their matrix clauses are processed in a way that mirrors the processing of coherence relations between independent clauses, causal RCs should be read faster than neutral RCs, which should in turn be read faster than concessive RCs. If, however, restrictive RCs are linguistic elements where language users do not typically expect to find information that is relevant at the discourse level, reading times should be fastest for neutral RCs since those RCs can be understood as simply disambiguating the referent; in contrast, reading times would be slowed by RCs whose content makes available a coherence relation and this pragmatic enrichment of meaning is posited to take time.

It has also been found that the content of causal relations is verified faster and recalled better than information from clauses that are not part of a causal relation (e.g., Sanders & Noordman 2000, Trabasso & van den Broek 1985, van den Broek 1990). In addition to comparing reading times of causal, neutral, and concessive RCs, Experiment 3 measured whether information provided by causal RCs is verified faster than information provided by neutral or concessive RCs. The verification statements only inquire about the contents of individual clauses, not about any discourse-level inferences, to avoid influencing participants' reading behavior as much as possible. Since the resulting statements are fairly easy to verify, we only measure reaction times, not the accuracy of responses.

6.5.1 Participants

52 monolingual English speakers were recruited through Amazon Mechanical Turk (mean age 40.22, age range 25-63, 31 women). They participated in exchange for monetary compensation (\$4.50).

6.5.2 Materials

Stimuli contained an introductory sentence that introduced or invoked a set of people from which one person would later be singled out by the restrictive RC construction, a target sentence consisting of a matrix clause with a direct object modified by an RC, and a wrap-up sentence. The target sentences differed in the coherence relation that could be inferred between the RC and the matrix clause (causal, neutral or concessive). The subject of all stimuli was a proper name or a first person pronoun,

while the direct object was a general NP that specified or implied a different gender or person than the subject; the subject was male in 50% of the items, female in the other 50%. The verb in the matrix clause was always an NP2-biased IC verb in the past tense. Each IC verb was matched with another IC verb to create (context-dependent) antonyms, e.g., *admire* and *pity*, or *thank* and *sue*. By manipulating the IC verb to change the coherence relation between the matrix clause and the RC, the RC was kept constant between conditions, see Table 1. Each set of IC verb antonyms was supplemented with a non-IC verb to create a neutral condition with the same RC. The full list of target items can be found in Appendix C.

Table 1
Sample item with target sentence in all three conditions

Intro	Jenny walked through the hallway to check on the daily goings-on around the office.
<i>neutral RC</i>	She joked with the guy who made a lot of money for the company.
<i>causal RC</i>	She praised the guy who made a lot of money for the company.
<i>concessive RC</i>	She fired the guy who made a lot of money for the company.
Wrap-up	She arrived at the conference room just in time for her next meeting.
Verification statement	<i>The guy made a lot of money for the company.</i>

The target items were distributed over three lists, with each item occurring only once per list, in one of the three conditions. Target items were interspersed with 12 ‘distractor’ fillers that also contained RCs, and 24 additional fillers of various types. Each participant saw every item only once, in one of the conditions.

Each item was accompanied by a verification statement. For the target items, the verification statement inquired only about the content of the RC and were always true. For the ‘distractor’ fillers that also contained RCs, the statements were all false and asked about various parts of the items. For the additional fillers, the statements were a mix of true and false, and asked about various parts of the stimuli. In total, a third of all verification statements were false; two thirds were true (see also Appendix C).

6.5.3 Procedure

Participants were recruited via Amazon Mechanical Turk, after which they were sent to another website, hosted by IbexFarm,³ where they could access the moving window self-paced reading experiment. Items were initially displayed as a series of horizontal lines on the screen; the length of the lines corresponded to the length of the regions. By pressing the space bar on their keyboard, participants could reveal the next region of the item. Items were presented non-cumulatively; when a new region was revealed, the previous region was again replaced by lines.

All target sentences were split up into two regions, with the matrix clause and the relative pronoun in the first region, and the rest of the sentence in the second region. The first and last sentences of every item were also presented over two regions. Each target sentence started on a new line and was followed by the first region of the wrap-up sentence. (20) illustrates the spatial configuration of target stimuli on the screen, with slashes indicating regions.

(20)

Jenny walked through the hallway to check on // the daily goings-on around the office //
She praised the guy who // made a lot of money for the company. // She arrived at the conference room //
just in time for her next meeting.

When finished reading the item, participants had to press the space bar once more to move on to the verification statement. They responded to the statement by clicking either TRUE or FALSE with their cursor. At six random moments in the experiment, participants were presented with a picture of a landscape. These pictures allowed participants to take a short break without it affecting the reading time measures. When they were ready to continue, participants clicked a “proceed” button at the bottom of the screen.

6.5.4 Data clean-up and analysis

For the analysis, we used residual reading times. Residual reading times were calculated using a regression equation that predicts the reading time of a region based on a participant’s reading speed and the length of the region; the predicted reading time is then subtracted from the actual reading time of the region (Trueswell, Tanenhaus, & Garnsey 1994). Residual reading times thus adjust for differences in the length of regions as well as differences in participants’ reading rates. Negative residual reading times indicate that a region was read faster than predicted, positive residual reading times indicate that a region was read slower than predicted. We

³ <http://spellout.net/ibexfarm/>

removed residual reading times that were more than three standard deviations above or below the mean (0.52% of the data).

All participants had a verification statement accuracy of above chance. As expected, the average percentage of correct responses was very high (93.94%), with 96.15% accurate responses to the target items and 92.47% to the filler items. The accuracy of responses per subject ranged between 76.67% and 100%. The reading time analysis was performed on all non-outlier data, regardless of whether the participant answered the item's verification statement correctly. The analysis of the reaction time to the verification statements was performed on correct responses only.

6.5.5 Results

6.5.5.1 Reading times

Table 2 provides an overview of the raw reading times per condition on the matrix clause, the RC, and the first region of the wrap-up sentence, which we have labeled the wrap-up sentence. Figure 4 shows the residual reading times on each of these regions for all three conditions.

Table 2
Mean raw reading times and standard deviations per condition per region in milliseconds

	<i>Matrix clause</i>		<i>RC</i>		<i>Spill-over</i>	
	M	SD	M	SD	M	SD
neutral RC	1243	667	1573	903	1302	763
causal RC	1207	695	1500	918	1271	702
concessive RC	1189	547	1706	1149	1350	813

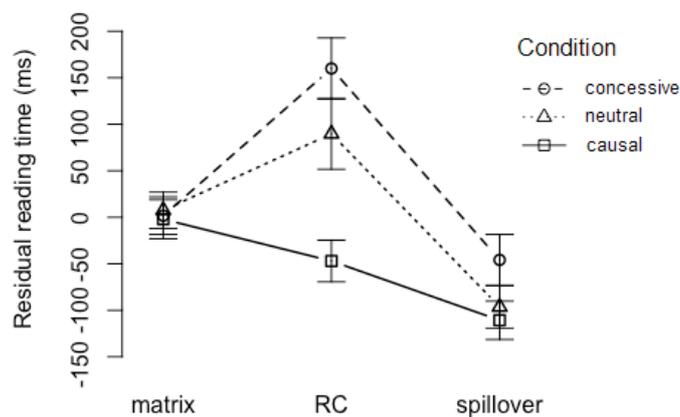


Figure 4. Residual reading times on the matrix clause, the RC, and the spill-over region per condition in milliseconds

We analyzed the residual reading times on the RC region in a linear mixed effects model. The best model contained both a main effect of condition ($\chi^2(1)=9.56, p<.01$) and a main effect of trial number ($\chi^2(1)=19.92, p<.001$); the effect of trial number indicates that reading times sped up as the experiment progressed. Pairwise comparisons of the condition variable revealed that causal RCs were read faster than concessive RCs ($\beta=-202.77, SE=37.96, t=5.35, p<.001$) and neutral RCs ($\beta=131.32, SE=39.61, t=3.32, p<.01$). Neutral RCs were only read marginally faster than concessive RCs ($\beta=-69.23, SE=41.55, t=-1.67, p=.08$). We found no main effects of condition on either the matrix clause ($\chi^2(1)=0.14, p=.70$) or the spill-over region ($\chi^2(1)=0.34, p=.56$).

6.5.5.2 Reaction times on verification statements

Only accurate responses were included in the analysis of the reaction times of the verification statements (96.15% of the data; 97.12% for causal RCs, 95.67% for concessive and neutral RCs). Table 3 includes the mean reaction times to the verification statements per condition. A likelihood ratio test revealed no main effect of condition ($\chi^2(1)=0.90, p=.64$).

Table 3

Mean reaction times and standard deviations of verification statements per condition in milliseconds

	M	SD
neutral RC	2202	946
causal RC	2205	923
concessive RC	2250	998

6.5.6 Discussion

The results of the self-paced reading experiment reveal that causal RCs are read faster than neutral and concessive RCs, which is in line with the relative processing ease of causal, additive, and concessive relations that hold between independent clauses. However, we did not find that concessive RCs were read slowest, which is what would be expected on the basis of evidence from earlier studies that compared reading times between negative and positive coherence relations.

Concessive relations are more complex and less expected by language users than causal or additive relations (see Chapter 4). In addition, they are much less often expressed without an overt linguistic marker (Chapter 4; see also Asr & Demberg 2012, Taboada 2006). Fine, Jaeger, Farmer, and Qian (2013:2) formulate the “rapid expectation adaptation” account, which states that “comprehenders are able to rapidly adapt to the statistics of novel linguistic environments.” In two self-paced reading tasks, Fine et al. (2013) show that processing disadvantages of linguistic constructions

that are usually rare diminish or even disappear if participants are repeatedly exposed to it in an experimental setting. We explored our data to check for a similar effect. We divided the dataset into the first and second half of the experiment, and plotted the residual reading times for both halves, see Figure 5. It thus seems that concessive RCs, the rarest construction in our experiment, was initially read slower than both causal and neutral RCs. As the experiment progressed, the difference in reading time between neutral and concessive RCs strongly diminished.

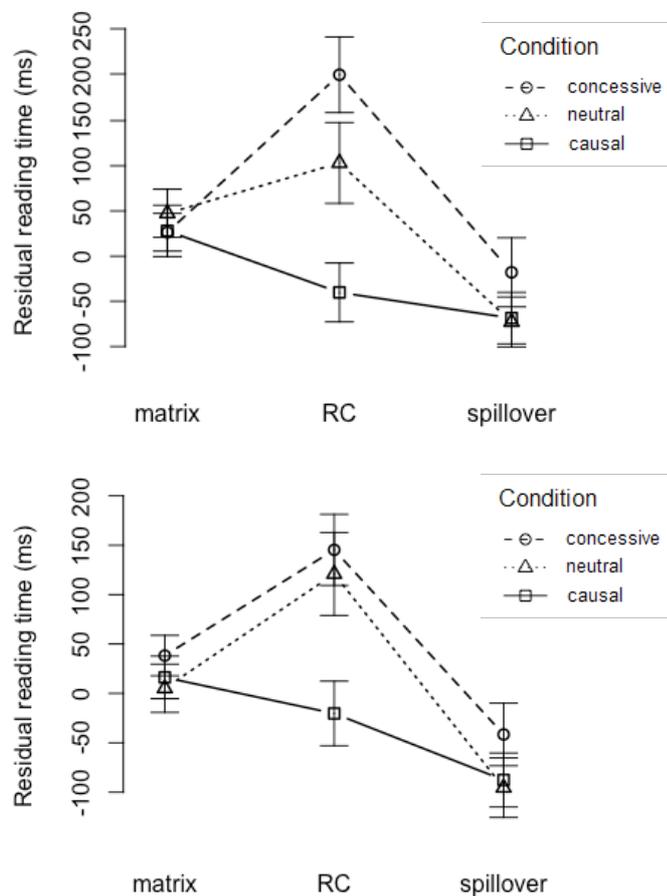


Figure 5. Residual reading times on the matrix clause, the RC, and the spill-over region per condition in milliseconds in the first (above) and second half (below) of the experiment.

Overall, the reading time results are more in line with processing studies on coherence relations between independent clauses than with reading time patterns that

would be predicted if language users do not expect restrictive RCs to contain information that is relevant at the discourse level. This suggests that readers naturally infer coherence relations between restrictive RCs and their matrix clauses, and that they do not only start making these inferences when the linguistic context (i.e., syntactic disambiguation) or the experimental setting encourages them to do so.

We found no differences between the conditions in the speed with which the statements were verified. This may be due to the way in which participants had to respond to the statements. After they finished reading the items, they pressed the spacebar one more time to move to the verification statement. They then had to respond to the statements by clicking on either the TRUE or FALSE button with their cursor. It was not possible in the web-hosted experiment to control the position of the participants' cursor as they entered the verification screen. This, in combination with the time it took participants to switch from the keyboard to the mouse, probably resulted in a fairly noisy measurement, since the cursor may have been anywhere, and plausibly at the edge of the screen so that it would not cover part of the self-paced reading items (participants were instructed to move their cursors out of the way). It is very possible that results would have been different if reaction time had, for instance, been measured using a button box in a lab set-up, since a button box would allow for a much faster and more consistent baseline response.

6.6 Experiment 4: Eye-tracking-while-reading

Experiment 3 focused on the processing of restrictive RCs by readers. Experiment 4 uses an eye-tracking-while-reading paradigm to investigate how restrictive RCs affect discourse-level expectations in on-line processing; specifically, how they affect expectations about upcoming coherence relations and expectations about upcoming referents.

In an off-line study, Kehler and Rohde (2015) show that IC verbs with a restrictive RC attached to the object receive fewer explanation continuations if the RC already provides an explanation for the matrix clause event (approx. 40% explanations) than if the RC merely provides additional information about the object (approx. 70% explanations). Experiments 1 and 2 in the current chapter show that the proportion of NP2 continuations following *because* is lower after a causal RC (approx. 50%) than after a neutral RC (approx. 80%). Both Kehler and Rohde's (2015) findings and the findings from Experiments 1 and 2 can be explained by the IC bias already being fulfilled by the explanation provided in the RC. Processing studies featuring IC verbs suggest that expectations on the basis of IC verbs are generated quickly (e.g., Koornneef & Sanders 2013, Koornneef & van Berkum 2006, Mak & Sanders 2013). This eye-tracking experiment examines whether and, if so, how fast language users update their discourse model with the interpretation that such discourse-level expectations have been met. If an RC provides an explanation for the

event denoted by the IC verb, *because* should be less expected than in a context where an RC merely provides additional information about the object of the matrix clause. Similarly, subjects that refer to the NP2 should be less expected in contexts where an RC has provided an explanation featuring the NP2 than in contexts where the IC verb event has not yet received an explanation.

Implicit causality verbs raise the expectation of an explanation relation (Kehler et al. 2008). Many implicit causality verbs have also been shown to be implicit consequentiality verbs; in result relations, i.e., forward causal relations, they also tend to display a next-mention bias, although this bias need not be toward the same referent as the verb's next-mention bias in explanation relations (e.g., Commandeur 2010, Crinean & Garnham 2006, Pickering & Majid 2007, Stewart, Pickering, & Sanford 1998). The expectation of a consequence after an IC verb, however, is not as strong as the expectation of an explanation, as is for instance shown by Experiment 3 in Kehler et al. (2008). In this experiment, we also test whether the expectation for a consequence increases when an explanation for an event denoted by an IC verb has been provided.

6.6.1 Participants

79 native speakers of English were recruited at Lancaster University. They participated in exchange for course credits. Data from four participants had to be discarded because of problems with the computer or eye-tracker. The data from the remaining 75 participants were analyzed (mean age 20.19, age range 18-41, 58 women).

6.6.2 Materials

Each participant was presented with 32 experimental stimuli, intermixed with 78 filler items. The filler items were taken from two unrelated experiments that consisted of stimuli that were similar to the target items in terms of length and complexity. The target stimuli consisted of an introductory sentence that introduced or invoked a set of people from which one person would later be singled out by the restrictive RC construction, a target sentence, and a wrap-up sentence. Target sentences consisted of a main clause with an NP2 IC verb of which the object NP was modified with a restrictive RC, a connective (*because* or *and so*), an adverbial, and a pronoun that referred to the NP2. The content of the RCs was designed to support the inference of different relations between the RC and main clause: neutral or causal. Table 4 contains a sample item in all conditions; the full list of materials can be found in Appendix D. We used *and so* instead of just *so* in the consequence condition to create a connective region that was approximately as long as *because* and to avoid the potential interpretation of *so* as meaning *so that*. To be able to measure reading times on the connective as well as on the pronoun, we added an adverbial, *clearly* in the

example in Table 4, to create a spill-over region for the connective. The region immediately following the pronoun was also kept consistent in each condition.

Table 4
Sample item with target sentence in all four conditions

Intro	Geoff was on a red-eye flight to New York.
<i>neutral RC</i> + <i>because</i>	He ridiculed the stewardess who was walking down the aisle because clearly she refused to acknowledge that she needed a dress in a much larger size.
<i>neutral RC</i> + <i>and so</i>	He ridiculed the stewardess who was walking down the aisle and so clearly she refused to provide him with any kind of service for the rest of the flight.
<i>causal RC</i> + <i>because</i>	He ridiculed the stewardess who crashed the drink cart into one of the seats because clearly she refused to acknowledge that she needed a dress in a much larger size.
<i>causal RC</i> + <i>and so</i>	He ridiculed the stewardess who crashed the drink cart into one of the seats and so clearly she refused to provide him with any kind of service for the rest of the flight.
Wrap-up	The other passengers thought Geoff was incredibly rude.

The target items were distributed over four lists. Each participant saw every item only once, in one of the four conditions. It should be noted that all IC verbs are NP2 verbs in the *because* conditions, but the next-mention bias for our set of IC verbs is mixed in the *and so* conditions. Our design is therefore 2x2 on the connective regions (connective and spillover), but we only have 2 conditions (*neutral+because* and *causal+because*) on the next-mention regions (pronoun and pronoun spillover), since NP2 is not the baseline next-mention bias for all of the IC verbs in the consequence condition.

6.6.3 Procedure

Participants were seated approximately 60 cm away from the monitor on which the experimental items were presented, with their heads on a chin rest. Eye movements were recorded using the SR Research Eyelink 1000 at the sampling rate of 500 Hz. Since the experiment was fairly long, it was split up over two sessions. After reading half of the experimental items, participants had a short break during which they performed another task, which did not involve a computer. After the break, participants returned to the monitor, were recalibrated, and finished the experiment. On average, the whole session took approximately an hour.

Participants were presented with a verification statement after 25% of all items (target or filler). The verification statements were included in the experiment to promote careful reading; no reaction time was measured. The verification statements

for the target stimuli always inquired about the context of either the first or last sentence of the item.

6.6.4 Data clean-up and analysis

Any fixations shorter than 80 ms and within one degree of a consecutive longer fixation were merged with the longer fixation. Any remaining reading times smaller than 80 ms were removed. Finally, outliers were removed in all reading times by replacing reading times of more than two standard deviations from both the participants' and the condition's mean by the value that corresponded, depending on the direction of the outlier, to either two standard deviations below or above the mean (2.0% of the data).

We analyzed four regions from the target sentences, see (21). We also performed an analysis in which we merged the pronoun and the pronoun spill-over region, since pronouns are very short and often not fixated on during reading, which leads to a lot of missing data points. Since this change did not impact our results, we only report the analysis with four regions.

(21) ... [because]_{conn} [clearly]_{connspill} [she]_{pron} [refused to]_{pronspill} ...

For each region, we analyzed three reading time measures: *first pass duration*, which is the time spent in a region until leaving the region for the first time in any direction; *regression path duration*, which is the time spent in a region plus all regressions to previous regions before leaving the region for the first time to the left; and *total fixation duration*, which is the total time spent in a region.

6.6.5 Results

Table 5 contains the mean reading times and standard deviations for each reading time measure per condition and region.

We did not find any significant differences at the connective region, the pronoun region or the pronoun spillover region ($p > .05$ for all measures). On the connective spill-over region, we found a main effect of connective in the first pass duration ($\chi^2(1)=11.58, p < .001$); *because* was read faster than *and so* ($\beta=38.78, SE=10.87, t=3.57, p < .001$). There was also a main effect of connective in the regression path duration ($\chi^2(1)=17.96, p < .001$); here too, *because* was read faster than *and so* ($\beta=89.78, SE=18.54, t=4.84, p < .001$). *Because* being read faster than *and so* is in keeping with the prediction that IC verbs generally favor subsequent explanations. In the total fixation duration, we found an interaction effect between connective and RC ($\chi^2(1)=7.96, p < .01$; this model does not include correlations between the random intercepts and random slopes for item and participant). This interaction is plotted in Figure 6. The impact of RC type on the total fixation duration is greater after *because* than after *and so*. This interaction pattern is in keeping with the prediction that causal

RCs can satisfy the expectation for an explanation and render the subsequent *because*-clause more surprising.

Table 5
Mean reading times and standard deviations per measure per condition per region in milliseconds

	Connective		Connective spill-over		Pronoun		Pronoun spill-over	
	M	SD	M	SD	M	SD	M	SD
First pass								
neutral + <i>because</i>	243	120	281	141	215	102	310	168
neutral + <i>and so</i>	254	146	322	211	218	94	344	212
causal + <i>because</i>	233	104	292	156	204	83	321	187
causal + <i>and so</i>	249	139	325	207	221	95	347	194
Regression path								
neutral + <i>because</i>	315	271	371	338	302	289	444	382
neutral + <i>and so</i>	320	294	480	430	327	371	496	380
causal + <i>because</i>	296	244	386	305	290	287	496	405
causal + <i>and so</i>	298	254	450	339	288	224	506	344
Total fixation								
neutral + <i>because</i>	334	217	393	254	260	154	450	304
neutral + <i>and so</i>	360	252	483	306	277	172	505	305
causal + <i>because</i>	340	200	422	281	253	142	474	297
causal + <i>and so</i>	344	217	466	293	290	169	506	287

All models for the connective spill-over region also included a main effect of trial number to account for participants' reading times speeding up throughout the experiment ($p < .05$ in all reported models).

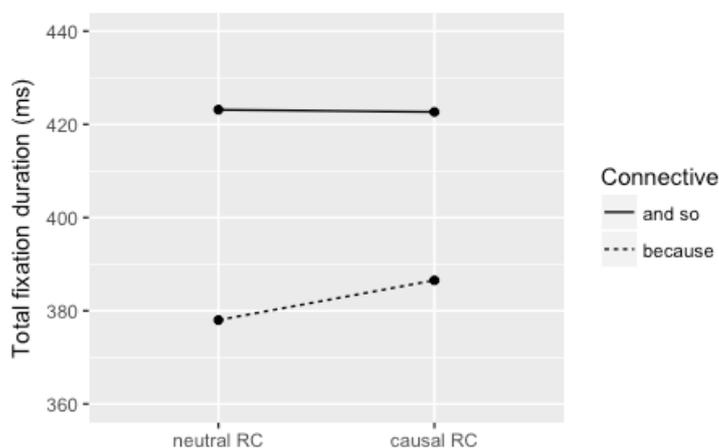


Figure 6. Total fixation duration at connective spill-over region

6.6.6 Discussion

The results suggest that readers update their expectations about upcoming explanation relations when processing RCs; a causal RC can reduce the expectation of an explanation. However, this effect only occurred in the connective spillover region, not on the connective itself. Furthermore, we only found an effect in the total fixation duration, not in measures more indicative of immediate processing difficulties (first pass gaze duration and regression path duration). The results thus suggest that while the effect of causal RCs on the expectation of explanation relations found in off-line studies (e.g., Kehler & Rohde 2015) can also be observed in on-line processing, the processing disadvantage of *because* after a causal RC as compared to after a neutral RC is not immediate.⁴ This suggests that participants were not instantly surprised or thrown off by *because* after an explanation. The effects not being immediate could be due to the fact that the coherence relation between the RC and its matrix clause is implicit and readers need to read the full clause before they are able to establish the causal coherence relation, at which point the next word is the word *because*. Alternatively, readers may take into account the possibility of *because* attaching to the RC, although the continuations from Experiment 2 suggest that causal relations within restrictive RCs embedded under an IC verb are not very expected.

Finally, the non-immediate effects could be due to the expectation of a coherence relation not being entirely the same as the expectation of a connective. Although coherence relations and connectives are closely related, there is no one-to-one relation between, for instance, explanation relations and *because*. Mak and Sanders (2013), for example, show that effects of IC verbs on processing still occur in explanation relations that are not marked by a causal connective, but by a temporal connective (*when*). Upon encountering the causal connective in their experimental items, the expectation of an explanation had been met, but if the IC verb also triggers the expectation for, or even primes, a causal connective, this would not necessarily be undone by an implicit explanation relation; the processing of *because* could still be (initially) facilitated because it was primed by the expectation of an explanation relation. The exact mechanism behind the updating of expectations about upcoming coherence relations thus seems worth further exploring in future research.

Our results did not support the hypothesis that a context where the IC verb expectation of an explanation relation has been met facilitates the processing of a consequence relation. Even though IC verbs have also been associated with consequence relations, we did not find any indication that consequences were more expected after a causal RC than after a neutral RC.

⁴ The effects found in the current eye-tracking experiment are less immediate than effects found in other eye-tracking studies involving IC verbs (e.g., Koornneef & Sanders 2013, Mak & Sanders 2013) and less immediate than effect found in more general eye-tracking experiments comparing the processing of a connective in different contexts (e.g., Canestrelli, Mak, & Sanders 2013, Kleijn 2012).

Finally, our eye-tracking data does not show the reduced expectation of an upcoming NP2 mention after a causal RC that we observed using off-line measures in Experiments 1 and 2. This could indicate that expectations about upcoming referents are updated less quickly than expectations about upcoming coherence relations. It could also be the case that the expectation of an NP2 mention is not reduced enough to lead to slower processing in the causal+*because* condition; Experiments 1 and 2 show that the proportion of NP2 mentions is still approximately 50% after a causal RC. Finally, the fact that our on-line processing data does not show the pronoun biases found in off-line studies could also reflect an asymmetry between pronoun comprehension and pronoun production, as observed by for instance Rohde and Kehler (2014); while eye-tracking-while-reading exclusively measures comprehension, continuation tasks involve both comprehension (of the prompt) and production (of the continuation). How exactly expectations about upcoming referents are updated during reading remains a question for future work. To further study the expectedness of pronouns after IC verbs, it seems useful to insert linguistic content between the connective and the pronoun, as in the stimuli used in the current experiment (if this is possible in the language in which the experiment is conducted); this makes it possible to differentiate between the processing of the connective and the processing of the pronoun.

6.7 General discussion and conclusion

Over the course of four experiments, this chapter explored whether coherence relations can be inferred between restrictive RCs and their matrix clauses. The results from all experiments indicate that language users treat restrictive RCs as linguistic elements that potentially contain information that is relevant at the discourse level, and that the contents of a restrictive RC can be linked to its matrix clause in a causal or concessive, i.e., negative causal, coherence relation. This implies that restrictive RCs should not be categorically excluded from receiving discourse segment status.

Although the results of all experiments indicate that language users can infer coherence relations between restrictive RCs and their matrix clauses, the influence of RCs on the expectations about upcoming referents found in Experiments 1 and 2 could not be observed in on-line processing in Experiment 4. Differences between results from off-line and on-line experiments that look into IC verbs have been observed before (see Koornneef & Sanders 2013:1172). A lot of psycholinguistic research has focused on whether language users have expectations about upcoming linguistic content, and how violations of those expectations are resolved, see e.g., Kutas, DeLong, & Smith (2011) and Kuperberg and Jaeger (2016) for overviews. Investigating how quickly expectations, or the effects of expectations, disappear after they have been met could be a promising topic for future research.

The idea that restrictive RCs cannot enter into a coherence relation with their matrix clauses or any other clauses is based on the fact that restrictive RCs are both syntactically and conceptually integrated in their matrix clauses. The RC is syntactically embedded in an NP and, as such, part of the matrix clause. In addition, the RC provides vital information about the referent it modifies, without which the matrix clause is conceptually incomplete. The experiments in this chapter suggest that neither of these factors make the restrictive RC completely unavailable as an independent discourse segment. In the causal and concessive experimental items, as well as in examples (2)-(6) given in the introduction, there appears to be a mismatch between the syntactic structure and the discourse structure of the sentence. In (22), for example, the RC relates to only the NP *the stewardess* at the syntactic level, while at the discourse level it relates to the entire matrix clause.

(22) Geoff ridiculed the stewardess who crashed the drink cart into the wall.

When it comes to the conceptual level, the RC appears to fulfill a double function. It restricts the referent so that, in the case of (22), it is clear which stewardess is ridiculed. At the same time, it provides a reason or explanation for the matrix clause event, in this case the ridiculing.

If restrictive RCs cannot be categorically excluded as discourse segments, how should they be treated in discourse segmentation and annotation? A liberal option would be to allow every restrictive RC to be a discourse segment. In the absence of a meaningful discourse-level relation between the RC and its matrix clause, the relation could be annotated as ADDITIVE (or using a similarly general relation label, such as ELABORATION or CONJUNCTION). The drawback of this approach is that the discourse structure and discourse annotations may include relations that are irrelevant at the discourse level; while the experiments in this chapter show that a restrictive RC *can* relate to its entire matrix clause at the discourse level, there is no reason to believe that this is *always* the case. A more conservative option would be to only segment a restrictive RC construction and annotate the coherence relation between the segments if the sentence allows for a discourse-level inference between the RC and its matrix clause. Since this segmentation procedure, much like segmenting the linguistic constructions discussed in Chapter 2, relies heavily on the interpretation of the annotator, it is not a very suitable option for automated segmentation or annotation applications. For manual segmentation and annotation procedures, however, this option would likely result in a more accurate representation of the discourse structure and a more complete overview of the coherence relations that hold in a discourse.

7 Conclusion and discussion

Understanding a discourse requires language users to infer coherence relations between segments in a text. This can be a fairly simple task if two clauses are conjoined by a connective that exactly specifies which relation should be inferred. Often, however, it is much more complex, as many of the examples throughout this dissertation illustrate. A connective may be absent or less specific than the relation that has to be inferred. It may also not be obvious which discourse segments should be related to each other, a problem especially relevant for contexts or sentences with multiple clauses. This dissertation focused on both discourse segmentation and the linguistic marking of coherence relations, specifically asking between which parts of a discourse language users infer coherence relations and when and why coherence relations are explicitly marked by a connective. In exploring both issues, this dissertation took a cognitive perspective. This final chapter first discusses the main findings of the project when it comes to discourse segmentation and the linguistic marking of coherence relation. In accordance with the goals of the overarching project of which this dissertation was a part, the findings are related to machine translation. Subsequently, this chapter discusses some outstanding questions and formulates topics for future research. Finally, it gives a brief overview of some issues that need to be addressed to further develop a cognitive theory of coherence relations.

7.1 Overview of main findings

7.1.1 Discourse segmentation

One of the questions this dissertation set out to explore was between which parts of a text language users infer coherence relations. This property of discourse coherence is captured through the practice of discourse segmentation. Chapter 2 argued that discourse segments should correspond to the idea units that are related to each other in the mental representation of the discourse. To accurately segment a text and to create the best representation of the discourse structure, it was argued, two well-established discourse segmentation guidelines should be loosened: the treatment of segmentation and annotation as a two-step process and the completeness constraint, which poses that all elements should be included in the segmentation of a text. Both guidelines were proposed by Mann and Thompson (1988) and have been widely used in annotation efforts and implemented in discourse annotation approaches ever since.

Separating annotation from segmentation makes the segmentation process more objective and prevents circularity (Taboada & Mann 2006). However, as Chapter 2 demonstrated, without interpretation it can be impossible to differentiate between fragments in which a relation holds between a matrix clause containing a complement and a subsequent clause, as in (1), or between two clauses embedded under a single matrix clause, as in (2).

- (1) [It makes me sad to know they'll never get to know each other]_{S₁}
because [my dad was a really wonderful person.]_{S₂}
- (2) [I was just telling my husband about how [I was so intimidated by a
new acquaintance]_{S_{1a}} because [she is so/too cool.]_{S_{1b}}]_{S₁}

Both (1) and (2) consist of a clause with a complement-taking predicate, another clause, and *because* followed by a third clause. However, as is indicated by the square brackets, the fragments differ in their discourse structure. The relations in (1) and (2) are representative of the larger issue addressed in Chapter 2, that fragments with identical surface forms can have very distinct underlying discourse structures. The ambiguity in (1) and (2) was created by relations potentially being syntactically embedded, but the same problem occurs with relations being embedded in other relations, as was annotated in Chapter 4. The sentences in (3) and (4) both contain three clauses and two connectives. While in (3), the first two clauses together make up the S₁ of the relation marked by *and*, *and* in (4) connects the second and third clause of the sentence (not counting the embedded clause that starts with *using*), which together make up the S₂ of the relation marked by the first connective *because*.

- (3) [[I carry pepper spray and a flash light with me]_{S1} when [I walk my dog at night (other times, too!)]_{S2}]_{S1} and [I always keep my wits about me.]_{S2}
- (4) [I don't recommend this method in this master recipe]_{S1} because [[most of our cast iron skillets have been heavily seasoned using savory foods like onions and garlic,]_{S1} and [those flavors can sometimes transfer to the sweet cake.]_{S2}]_{S2}

Chapter 2 thus showed that distinct discourse structures can result in the same surface structure, as illustrated by (1)-(2) and (3)-(4), which makes it necessary to incorporate interpretation into discourse segmentation; only then can we arrive at a discourse segmentation that accurately captures the idea units that are related to each other in the mental representation of a text.

Chapter 2 also discussed the completeness constraint. The completeness constraint poses that all elements of a text should be included in the segmentation of that text. In Chapter 2 it was argued that the completeness constraint should be amended to state that all elements of the propositional content of a text should be included in the segmentation of that text; the inclusion of other elements, for instance stance markers, should be optional. In (5), the stance marker *probably*, is not part of the propositional content of the sentence, but it is an integral part of S₂ of the relation signaled by *because*; it is the likelihood of someone already having ordered extra fried baby artichokes that leads to the claim that eating the last of the current artichokes is no problem. The function of *probably* in the relation in (6), on the other hand, is different. It is located in S₁, but appears to take scope over the entire relation; it is likely that if the toddler was allowed to eat 7 yogurts per day, he would do so. Including *probably* in S₁ of (6) therefore results in less accurate segments than including *probably* in S₂ of (5). A more appropriate option to segment (6) would be to remove *probably* from S₁ or otherwise indicate that it modifies the entire relation, even if that means that *probably* is not accounted for in the segmentation of the discourse. Note that this decision also involves interpretation.

- (5) [It's no big deal if someone eats the last fried baby artichoke]_{S1} because [someone has probably already flagged down the waiter and ordered more.]_{S2}
- (6) [He would probably eat 7 per day too]_{S1} if [he could!]_{S2}

Finally, Chapter 6 returned to the question of between which parts of a discourse language users infer coherence relations. Chapter 2 listed restrictive relative clauses (RCs) as being commonly excluded from receiving discourse segment status. This would, however, only be appropriate if language users indeed never infer a coherence relation between a restrictive RC and its matrix clause. The parallel corpus used as the basis for Chapters 4 and 5 included several examples of coherence

relations being translated with restrictive RC constructions. This suggests that perhaps restrictive RCs should not be categorically excluded as discourse segments.

Chapter 6 used a combination of experimental methods to explore the availability of discourse-level inferences between restrictive RCs and their matrix clauses. By means of two continuation experiments, a self-paced reading experiment, and an eye-tracking experiment, it was shown that language users can indeed infer a coherence relation between a restrictive RC and its matrix clause. The range of available coherence relations does not only include causal relations, as in (7), but also appears to extend to other types of coherence relations, such as relations with a negative value for POLARITY. In (8), there is a contrast, or even a denial of expectation between the matrix clause of the underlined clause and the restrictive RC; even though he was homeless, she fell in love with him.

- (7) A man who murdered his ex-partner's new girlfriend by cutting her throat at her workplace has been jailed for 26 years.
- (8) I was 30, single, with a successful career when I fell in love with a man who lived in a bush. A man with no income. Or career prospects. Or shoes.

Chapter 6 concluded that while restrictive RCs connect to a noun phrase at the level of syntax, sometimes they relate to the whole matrix clause at the discourse level. Restrictive RCs should therefore not be excluded from receiving discourse segment status. However, since there is no reason to assume that restrictive RC *always* relate to the whole matrix clause at the level of discourse, determining for each restrictive RC whether it should be segmented is likely to result in the most accurate representation of the discourse. This, much like the segmentation of (1)-(6), requires the interpretation of a fragment.

The insights of both Chapter 2 and 6 help create better and more accurate discourse segmentations, which is important for several reasons. First of all, discourse segmentation is a crucial step in the discourse annotation process, and variability in segmentation can have consequences for the relation labels attributed to relations during annotation (Demberg, Asr, & Scholman 2017). Incorporating interpretation into segmentation might make the process more circular, but it helps separate segmentation and annotation when it comes to inter-annotator agreement. Asking coders not just which relation holds between two segments, but also between which two segments a relation holds can help pinpoint sources of disagreement. This would make the annotation process as a whole more transparent. Discourse segmentation, however, is not solely instrumental to the process of discourse annotation. It also creates a representation of the hierarchical structure of a text, which is an aspect of discourse worth studying in itself. For instance, the finding that coders sometimes

need their interpretation of a fragment to determine the structure of a discourse suggests that language users, when processing language, may also need to relate a segment to the preceding discourse based on the relation that is most plausible given the content of the segments. Finally, discourse structure seems to influence the linguistic marking of coherence relations, both when it comes to *whether* relations are marked (explicit versus implicit) and when it comes to *how* relations are marked (connective choice); both of these points are touched upon in the next two sections.

7.1.2 The linguistic marking of coherence relations

The other main research question of this project focused on when and why coherence relations are explicitly marked by a connective. Coherence relations can be linguistically signaled by means of connectives (e.g., *but*, *because*) or cue phrases (e.g., *on the other hand*, *which is why*), but can also be left implicit and conveyed through the simple juxtaposition of two clauses or sentences. However, it seems that not all relations are equally easy to reconstruct when they are implicit. Removing the connective from (9), for instance, results in an implicit relation that is plausibly interpreted in the same way as the explicit version. Removing the connective from (10), on the other hand, makes it unlikely to be interpreted as a conditional relation.

- (9) [They would never get over that night]_{S1} because [they had said things that altered how they saw each other.]_{S2}
- (10) If [we pistol-whipped him like he said we did,]_{S1} [it wasn't me.]_{S2}

The intuition that some types of relations can be more easily left implicit than others is confirmed by analyses of discourse-annotated corpora (e.g., Asr & Demberg 2012, Taboada 2006). While the existence of asymmetries in the marking of coherence relations has been clearly established, the exact mechanisms that cause these asymmetries are not yet fully understood. This dissertation used parallel corpora to explore what influences whether or not a coherence relation is marked by a connective. Both Chapter 4 and 5 departed from the assumption that connectives are more likely to be used when they contribute essential information to the text; if a relation is already clear without a connective, it can easily be left unmarked. Chapter 4 investigated the influence of default expectations on the marking of coherence relations, while Chapter 5 focused on the presence of segment-internal cues that can function as signals for coherence relations. Both chapters used as their basis a discourse-annotated parallel corpus consisting of original English coherence relations and translations of those relations into Dutch, German, French, and Spanish. The English relations were annotated using the Cognitive approach to Coherence Relations (CCR). CCR was proposed by Sanders, Spooren, and Noordman (1992) and has been extended and refined in numerous other papers since; Chapter 3 provided an overview of state-of-the-art CCR for discourse annotation.

The translations of the English source text relations in the annotated parallel corpus were quantitatively analyzed in Chapter 4. The rate of implication of relations in translation was taken to be indicative of the ability of relations to be implicit in monolingual language use; if a specific type of relation can easily be left implicit, it will often be possible to leave out the connective in translation. The general hypothesis of the study was that expected relations are more often implicit than relations that are not expected, which is in line with predictions made by for instance the Uniform Information Density Hypothesis (Frank & Jaeger 2008, Levy & Jaeger 2007), effort and effect (key notions from Relevance Theory; Sperber & Wilson 1985, Wilson & Sperber 2005), and Horn's (1984) Q and R principles. More specifically, it was hypothesized that, by default, cognitively simple relations are more expected than relations that are cognitively more complex (see also Traxler, Bybee, & Pickering 1997). To determine the cognitive complexity of each relation, the relative complexity of all relevant CCR primitive values was determined using evidence from logic, language acquisition, language processing, and Mental Space Theory (Fauconnier 1985). The primitive values attributed to each English relation during annotations were then used to model the proportion of implicit versus explicit translations in the corpus. The results suggest that the implicit versus explicit marking of coherence relations is indeed influenced by cognitive complexity, with simple relations being more often implicit than relations that are more complex.

In addition to the cognitive complexity of relations, the corpus study in Chapter 4 took into account the position of relations in the discourse structure. Based on findings by Patterson and Kehler (2013), it was hypothesized that relations that contained another relation or shared a segment with another relation would have a higher likelihood of being implicitated, while relations that were embedded in another relation would have a higher likelihood of remaining explicit. Relations embedded under a syntactic construction were also predicted to have a higher likelihood of remaining explicit. While the results of the corpus study were largely in line with these hypotheses, there appeared to be a confound between relation type and relational and syntactic dependency. Specifically, there was a strong relation between the *BASIC OPERATION* and the dependency measures; the type of *BASIC OPERATION* that is least often implicit (conditional) is also the type of relation that is most often embedded under a syntactic construction or in another coherence relation. Conversely, the types of *BASIC OPERATION* that are most often implicit – causal and additive relations – also most often share a segment with or contain another relation. In the final analysis of Chapter 4, only containing another relation and sharing a segment with another relation were significant predictors of the implicit versus explicit marking of coherence relations, both promoting implicitness.

Chapter 5 focused on the contribution of non-connective elements to the marking of coherence relations. Traditionally, relations with a connective, such as (9)

or (10), have been labeled ‘explicit’ coherence relations; relations without a connective, as in (11), as ‘implicit’ relations. Although this distinction seems very straightforward, it is not without its problems. Connectives can for instance signal a relation that is less specific than the relation that is constructed by language users, as in (12), where the relation is marked by *after*, a temporal connective, but the inferred relation is causal. The relation in (12) is therefore less explicitly signaled than the relations in (9) or (10). In addition, a relation without a connective may contain strong other cues that help language users infer the appropriate relation. The semantic opposition in (13), for instance, could be argued to function as a signal for the contrastive coherence relation. This relation is then more explicitly signaled than the relation in (11), even though neither fragment contains a connective.

- (11) *I broke down and bought a Walkman* – [which surprised me.]_{S1} Ø [I had always ranked them between boa constrictors and Planet Hollywood T-shirts in terms of vulgar accessories.]_{S2}
- (12) *I started the day with a ceramic pig but* [abandoned it]_{S1} **after** [it got to be a drag to carry.]_{S2}
- (13) *The word therapy suggested a profound failure on my part.* [Mental patients had therapy.]_{S1} Ø [Normal people did not.]_{S2}

Connectives are the only linguistic elements that by definition express relational meaning, but that does not necessarily mean they are the only indicators for coherence relations. Chapter 5 explored how linguistic elements inside the discourse segment interact with connectives and how segment-internal elements contribute to the signaling of the coherence relation. Three distinct ways in which segment-internal elements can interact with connectives were formulated on the basis of existing literature: *division of labor*, *agreement*, and *general collocation*. The existence of this three-way distinction was demonstrated using data from the discourse-annotated parallel corpus. In division of labor types of interactions, the connective and the other signal overlap in the meaning they encode, and the presence of one is likely to make (part of) the other redundant. In this type of interaction, the segment-internal element functions as a signal for coherence relations because it explicitly signals (part of) the coherence relation. In agreement types of interactions, the connective and the other signal overlap in the meaning they encode, but they are commonly used in addition to each other. Because of the frequent co-occurrence of a segment-internal element and a connective in agreement types of interaction, the segment-internal element is predicted to raise the expectation for the relation. In general collocation types of interactions, there is no overlap in the meaning signaled by the connective and the other signal. In this type of interaction, the segment-internal element functions as a signal for a relation mainly because it frequently co-occurs with a specific type of relation and, as such, raises the expectation for that relation.

In general, relations that contain a segment-internal signal are hypothesized to less often contain a connective than relations that do not contain such a signal. However, not all non-connective signals for coherence relations appear to function as signals in the same way. The presence of a segment-internal signal is a much stronger predictor for the absence of the connective in division of labor types of interactions than in agreement or general collocation types of interactions. While the presence of a segment-internal signal may still increase the likelihood of more relations without a connective in agreement and general collocation types of interactions, the mechanism appears much weaker here and more susceptible to other factors influencing the explicit versus implicit marking of coherence relations.

Chapter 4 and 5 thus addressed three mechanisms that appear to influence whether or not a coherence relation is marked by a connective: default expectations, the position of the relation in the discourse structure, and the presence of other signals. It should be noted that these factors are by no means incompatible. Speakers or writers are for instance predicted to not mark a complex relation by a connective if another element in the relation already signals the relation, or mark a cognitively simple relation because it is embedded in another relation. Similarly, listeners or readers may by default expect a cognitively simple relation, but adjust this expectation upon encountering a segment-internal element that is strongly associated with a more complex type of relation. Table 1 gives an overview of the factors influencing the explicit versus implicit marking of coherence relations addressed in this dissertation. Factors that were empirically tested in the current project are marked in bold; factors that are hypothesized to influence the explicit versus implicit marking of coherence relations on the basis of the current project are given in regular font. The parallel corpus study in Chapter 4 found, much like Patterson and Kehler (2013), that embedded relations were more often implicit than non-embedded relations. However, since there turned out to be a confound between embeddedness and relation type, embeddedness was not a significant predictor in the final model. The exact role of embeddedness in the explicit versus implicit marking of coherence relations thus requires further investigation in future research.

Table 1
Overview of factors that influence the explicit vs. implicit marking of coherence relations.

Type		Increases the likelihood of
<i>Cognitive complexity</i>	Simple relation	Implicitness
	Complex relation	Explicitness
<i>Position in the discourse structure</i>	Includes another relation	Implicitness
	Shares segment with other relation	Implicitness
	Embedded in another relation	Explicitness
	Embedded under a syntactic construction	Explicitness
<i>Presence of segment-internal signal</i>	Overall	Implicitness
	Division of labor	Implicitness – strong

7.2 Discourse in machine translation

The research reported in this dissertation was conducted as part of a larger project focused on discourse in machine translation (MT), MODERN. The introduction illustrated several discourse-level problems in statistical MT pertaining to discourse segmentation, discourse structure, and the explicit marking of coherence relations: the use of incomplete connectives or cue phrases, wrong placement of a target text connective relative to the segments of the relation, inappropriate connective choice, and inappropriate implicature of the connective. This dissertation produced several insights and resources relevant to the improvement of MT at the discourse level.

First of all, the current project produced a discourse-annotated parallel corpus, which was used as a basis for Chapters 4 and 5. The corpus will be made publicly available (Hoek, Zufferey, Evers-Vermeul, & Sanders *in prep*). For MT purposes, the parallel corpus can be used as a practical resource for training an automatic sense labeler for connectives, which can be used as a step before automatic translation to improve the quality of connective choice using existing methods (Meyer, Hajlaoui, & Popescu-Belis 2015). In addition, the corpus provides an extensive overview of Dutch, German, French, and Spanish translation equivalents of eight highly frequent English connectives. The translation equivalents include both connectives and other linguistic elements. Since the manual translation spotting method used in creating the annotated corpus is much more reliable than automatic translation spotting (e.g., Cartoni et al. 2013), the list is bound to be more accurate and include more non-obvious translation equivalents than automatically generated resources. This assumption is supported by the implicature rate of connectives reported in Chapter 4 (4%), which is considerably lower than implicature rates found by automatic translation spotting systems, such as the 18% reported by Meyer and Webber (2013). The translation equivalents included in the corpus could thus be used as training data for MT systems or for automatic translation spotting systems used to generate annotated training data for MT systems.

At a more theoretical level, the current project contributed new insights about the implicitation of connectives in translation. This knowledge could for instance serve to post-process MT output; if, for instance, no translation equivalent of a conditional connective can be located in the MT output, a linguistic marker signaling conditionality likely has to be added. A similar correction may not, however, be necessary for positive causal or additive relations. In addition, the findings from Chapter 4 indicate that the implicitation of coherence relations also depends on the position of the coherence relation in the discourse structure; relations that contain another relation are for instance more likely to be implicitated in human translation. Producing human-like MT output may therefore require information about discourse structure, a variable that thus far is not commonly included in MT systems.

Another indication that information about discourse structure may eventually be crucial for automatically generating accurate and idiomatic translations is the existence of coherence relations embedded under syntactic constructions or under the scope of stance markers, as was discussed in Chapter 2. Example (7) from the Introduction, repeated here as (14), illustrates the problem this phenomenon poses to MT.

- (14) EN Can you imagine [all automotive plants being given the possibility of compensation]_{S1} **because** [they are located in the peripheral regions in the European Union]_{S2}? {ep-97-05-13}
- NL Kunt u zich voorstellen dat [alle autofabrieken de mogelijkheid krijgen tot compensatie]_{S1} **omdat** [ze gelegen zijn in perifere gebieden van de Europese Unie]_{S2}?
- MT [Kunt u zich voorstellen dat alle installaties in de automobielinindustrie krijgen de mogelijkheid van compensatie,]_{S1} **want** [ze zijn gevestigd in de perifere regio's in de Europese Unie?]_{S2}
'Can you imagine all automotive plants being given the possibility of compensation, since they are located in the peripheral regions in the European Union?'

As is indicated by the segmentation in the English original, the relation marked by *because* holds within the complement of *can you imagine*. The human Dutch translation preserves the discourse structure of the fragment and, as a result, the overall meaning of (14). The MT output, however, uses *want* 'because' to translate *because*. However, *want*, unlike *omdat* used in the human translation, cannot be embedded. The brackets in the MT version of the fragment indicate the change in discourse structure; the interpretation of the relation changes from an objective causal relation embedded in a question to a speech act relation in which a question is asked and motivated. For translation, it is thus crucial that embedded relations are translated

using a connective that can be embedded (see also Hoek, Evers-Vermeul, & Sanders 2016).

Including automatic text segmentation in an MT system may not yet be a very attractive option at this point because of the error rates of automatic text segmentation and the processing cost of such a segmentation step. A possible short-term solution would be to locate, in the source text, any specific contexts that are especially prone to ambiguities in discourse structure and manually check the translations of these fragments. Several linguistic elements that can give rise to discourse structural ambiguities were identified in Chapter 2: complement-taking predicates, relative clauses, and stance markers.

At the beginning of the MODERN project in 2014, statistical MT was the state of the art. Since then, neural MT has become increasingly popular and, in general, generates better translations than SMT (e.g., Isabelle, Cherry, & Foster 2017, Junczys-Dowmunt, Dwojak, & Hoang 2016, Koehn & Knowles 2017). Unlike SMT, NMT does not use phrases or n-grams, but models full sentences within a larger model. It could thus be expected that NMT could be trained to be more sensitive to discourse-level features than SMT, although, to the best of our knowledge, no research has yet evaluated the performance of NMT when it comes to the translation of coherence relations. However, it does seem likely that mistakes like the one in (4) in the Introduction, repeated below as (15), can be expected to be much less frequent in NMT than in SMT. In this example, the SMT output only includes the first half of a two-part ‘compound connective’ (*weliswaar ... maar*). As was explained in the introduction, this mistake is likely caused by the fact that the two parts of the connective are usually further apart than local context considered by SMT, while an attention-based NMT model could learn to cope with the long distance between *weliswaar* and *maar* (see for instance Miculicich Werlen, Pappas, Ram, & Popescu-Belis 2018).¹ For similar reasons, it could be expected that NMT makes fewer mistakes when it comes to the positioning of the connective relative to the two segments, as in example (16), (5) from the Introduction.

¹ NMT for instance greatly outperforms SMT when it comes to word order, especially when major changes have to be made between source text and target text (Bentivogli, Bisazza, Cettolo, & Federico 2016). This indicates that NMT has the potential to do much better than SMT when it comes to modeling longer-distance dependencies.

- (15) EN **Although** [the European Union has been very active in the nuclear sector,]_{S1} [up to now there has been a clear failure of the European Union and G-7 policy.]_{S2} {ep-99-03-10}
- NL **Alhoewel** de Europese Unie zeer actief is in de nucleaire sector heeft het beleid van de Europese Unie en de G-7 tot nog toe volledig gefaald.
- MT De Europese Unie heeft **weliswaar** zeer actief geweest in de nucleaire sector **Ø** tot nu toe is er sprake van een duidelijke falen van de Europese Unie en de G-7 beleid.
'The European Union has WELISWAAR been very active in the nuclear sector, Ø up to now there has been a clear failure of the European Union and G-7 policy.'
- (16) EN **Although** [it would have been better if we had voted at lunchtime,]_{S1} [I am happy to go ahead now.]_{S2} {ep-00-03-16}
- NL **Ofschoon** we deze stemming beter hadden gehouden rond de middagpauze, vind ik het goed als we het nu doen.
- MT **Maar** het zou beter zijn geweest als we hadden gestemd tegen de middag, ben ik blij om verder te gaan.
'But it would have been better if we had voted at lunchtime, I am happy to go ahead now.'

When it comes to other types of mistakes in the translation of coherence relations, there do not seem to be clear indications for why they should no longer occur with NMT. The main findings outlined in this section are therefore predicted to also be relevant to the improvement of NMT.

7.3 Topics for discussion and future research

Each chapter in this dissertation contained a discussion section that, if relevant, reflected on unanswered questions and limitations of the approach, and suggested areas for further research. This section discusses some outstanding research questions and potential topics for future research, focusing on issues that exceed the individual chapters.

7.3.1 Parallel corpora

Chapters 4 and 5 both used parallel corpora, which consist of original fragments and translations. As was elaborated on in the Introduction, translations can be a valuable tool in researching 'meaning,' since they can provide insight into the contribution of individual linguistic elements or constructions to the overall meaning of a sentence or text fragment. The current project used directional parallel corpora extracted from the Europarl corpus (Koehn 2005) to investigate the implicit versus explicit marking of coherence relations and the contribution of segment-internal elements to the marking

of coherence relations in monolingual language use; through the use of multiple language pairs, both issues could be studied cross-linguistically. At the same time, the parallel corpora were used to gain insights into the translation of coherence relations by human translators. Other applications of parallel corpora within the field of discourse can for instance be found in Cartoni, Zufferey, and Meyer (2013), Cartoni, Zufferey, Meyer, and Popescu-Belis (2011), Hansen-Schirra, Neumann, and Steiner (2007), and Levshina and Degand (2017).

The current project made use of the Europarl corpus mainly for practical reasons; it is the largest available parallel corpus with many language pairs. The main drawback of using Europarl seems to be that the language used in the European Parliament is not entirely representative of common language use; it is neither entirely spontaneous nor entirely prepared, it is spoken but also partly read out, it is at times highly formulaic, and sentences are often long and complex – it should be noted that throughout the dissertation, shorter examples were chosen if possible for the sake of clarity and word count. In the end, however, the benefits of Europarl were decided to outweigh the drawbacks. A later attempt to find or create a parallel corpus on which the corpus study reported in Chapter 4 could be replicated (to check for potential influences of genre) once more confirmed the limited availability of parallel corpora for the current research purposes. Other existing corpora either did not include all language pairs studied in Chapter 4 (UN corpus; Ziemski, Junczys-Dowmunt, & Pouliquen 2016), were too small, or were not very suited for studying discourse-level phenomena (OpenSubtitles; Lison & Tiedemann 2016 <http://www.opensubtitles.org/>; subtitling often involves shortening and simplifying utterances). Creating a new parallel corpus within the scope of the current project also proved unfeasible, either because of an insufficient availability of data or because of copyright, monetary, and time-related reasons.

Since the findings of the current project could not be checked against another parallel corpus, the extent to which the Europarl genre influenced the results cannot be answered definitively. It is expected, however, that for the ways in which translation data are used in this dissertation, differences between parallel corpora are mostly to be found in absolute, rather than in relative frequencies, and that the conclusions drawn on the basis of the translation data would not be different if another corpus had been used. Other parallel corpora may for instance contain a higher overall implicitation rate – either due to the translation process or to a higher level of redundancy in the source text, for instance because the register is highly spontaneous and in no way edited or pre-planned – or the frequency of specific relation types in the source text may differ between parallel corpora. It would still be expected that cognitively simple relations are implicated in translation more often than relations that are cognitively more complex, and that translators make use of similar target text equivalents to translate the source text relations (even though here absolute frequencies may differ as well). The main consequence of different implicitation rates

between parallel corpora for the studies reported in this dissertation would then concern the amount of data needed to reach enough observations to analyze.

While the Europarl corpus was a sufficient and invaluable resource for this project, it seems imperative to invest in creating new parallel corpora or extending existing parallel corpora in the future. Currently, the availability of large-scale parallel corpora with a wide variety of language pairs is highly limited, which is regrettable considering the wide-ranged and promising opportunities parallel corpora present for linguistic research.

7.3.2 The relationship between hierarchical discourse structure and relation type

In studying implicit versus explicit coherence relations, this dissertation mainly operationalized ‘implicitness’ as ‘unexpressed meaning.’ The main analysis of Chapter 4, for instance, categorized translations by means of a less specific connective than the ST connective as implicatures; when *because* is translated with a temporal connective, the causality is no longer explicitly encoded by the connective. Similarly, underspecified connectives in general were argued to be more implicit than fully specific connectives in Chapter 5. There are, however, indications that there is also a relationship between the syntactic configuration of a relation and the types of coherence relations that can be inferred. For instance, a relation in which the two segments are connected by *and*, a highly general connective that allows for the inference of much more specific types of relations, does not seem to be able to express all relations that can be expressed using the juxtaposition of two segments, and vice versa (Carston 2002, Crible & Demberg 2018); non-basic causal relations (prototypically signaled by *because*), for example, can be expressed by juxtaposed sentences, but not by a construction in which two clauses are connected by *and*. Similarly, the range of coherence relations that can be inferred between a main clause and a free adjunct is not equal to the range of relations inferable between two juxtaposed sentences (Kortmann 1991, Reid 2016); for example, juxtaposed sentences can express relations with a negative value for POLARITY, but free adjunct constructions cannot. In the absence of a connective or in contexts with an underspecified connective the syntactic structure thus may impose constraints on the range of available coherence relations; as such, the syntactic configuration of a relation could potentially be considered to be an explicit linguistic signal of coherence relations as well.

Chapter 4 took into account the position of a relation in the discourse structure in examining the factors that influence whether or not a relation is explicitly marked by a connective. One of the findings of the corpus study was that there appeared to be a relationship between relation type and the position of the relation in the discourse structure. Conditional relations, for instance, were more often found to be embedded

under a syntactic construction or in another relation than other types of relations. This finding suggests that the way in which different types of coherence relations tend to be embedded in the rest of the discourse is in some way structured. The nature of the relationship between relation type and discourse structure, both when it comes to the ordering of relations in the discourse structure and the syntactic configuration in which relations appear, would be worthwhile to explore further in future research.

7.3.3 Expectations about hierarchical discourse structure

Chapter 2 discussed the issue of ambiguity in discourse segmentation and identified linguistic constructions that seem especially prone to discourse structural ambiguity. When determining which segments are related to each other, coders at times need to take into account their interpretation of the fragment to resolve ambiguity; this suggests that this might also be the case for language users. A valuable step for future research could be to examine how language users deal with discourse structural ambiguities, for instance by investigating the relative frequencies of distinct discourse structural configurations in specific contexts and relate this to processing behavior. For example, embedded relations may be much more frequent after cognitive verbs (*think, believe*) than after other types of complement-taking predicates, which would lead to the prediction that embedded relations are more easily processed in contexts following a cognitive verb. Study 2 from Chapter 6 gave some insight into how willing language users are to include coherence relations inside restrictive relative clauses – in general, they did not seem too eager to construct relations within restrictive RCs. However, this was only in a very specific context with a main clause element that has shown to be highly salient at the discourse level (IC verb) and does not compare low attachment in RCs to low attachment in other types of constructions.

In addition, future research may focus on the more general questions of how sensitive people are to ambiguity in discourse segmentation and whether and to what extent people make predictions about discourse segmentation while processing language. After all, it is impossible for language users to spend as much time on determining which segments are related to each other when processing language as researchers do while analyzing a text. This makes it plausible that, much like seems to be the case for establishing the *type* of coherence relations that hold in a discourse, people have certain processing strategies or default expectations about discourse segmentation.

7.4 A cognitive theory of coherence relations

This project considered coherence relations to be cognitive constructs, using the Cognitive approach to Coherence Relations (CCR) as the basis for depicting and classifying coherence relations. While CCR was mainly used instrumentally throughout the dissertation, working within the CCR framework and applying its

primitives led to several new insights about the approach. First of all, Chapters 4 and 5 corroborate the claim made in the original Sanders et al. (1992) proposal that the CCR taxonomy provides a fruitful basis for explaining the phenomenon of implicit coherence relations (see also Spooren 1997 for an account of underspecified coherence relations rooted in the CCR approach). In addition to CCR's cognitive primitives making meaningful predictions about language acquisition and language processing (see Sanders & Evers-Vermeul *in press* for an overview), CCR's primitives thus also make meaningful predictions about the explicit versus implicit marking of coherence relations. In addition, the cognitive status of CCR's primitives implies that they refer to concepts beyond coherence relations and connectives as well. Stukker, Sanders, and Verhagen (2008), for instance, show that distinctions made in CCR can also be relevant to verbs. Chapter 5 of this dissertation showed that CCR's primitives cannot only capture what is signaled by connectives, but can also be used to depict what is signaled by other linguistic elements within coherence relations.

One of the appeals of using CCR to describe coherence relations is that the primitive value combinations reflect how closely relations resemble each other ('relations among relations'); positive subjective causal relations with a non-basic order are for instance similar to positive subjective causal relations with a basic order. It was observed in Chapter 3, however, that it might not *always* be the case that relations that differ in one value are more similar to each other than relations that have different values for two or more distinctions. Specifically, it was observed that negative causal objective relations and negative additive subjective relations are often confused in annotation. The difficulty of distinguishing between negative causal and negative additive relations (often referred to as concessive relations and contrast relations) has also been noted in other annotation projects (e.g., Robaldo & Miltsakaki 2014, Zufferey & Degand 2017). While the cognitive relevance of each of CCR's primitives has been tested or demonstrated in many different studies (Sanders & Evers-Vermeul *in press* for an overview), it may thus be valuable to investigate the cognitive plausibility of *combinations* of primitives. It could for instance be the case that the BASIC OPERATION primitive is more important or salient in the domain of positive coherence relations than in the domain of negative coherence relations. Testing to what extent language users can differentiate between negative additive and negative causal relations in off-line tasks (e.g., annotation, paraphrase task, or a task in which participants group together relations that are similar to each other) and whether these types of relations display distinct patterns in on-line language processing experiments can help verify whether all value combinations of POLARITY and BASIC OPERATION are cognitively plausible as well.

Chapter 3 gave an overview of proposals for new distinctions within the CCR framework since the original 1992 taxonomy. The chapter mainly focused on CCR as

a basis for discourse annotation; all of the new distinctions contribute to the descriptive adequacy of CCR. For distinctions to be adopted into CCR, they have to meet several criteria. They have to be properties of the relational surplus (relational criterion), they have to describe relations that hold between clauses or larger discourse segments, and they have to be cognitively plausible. Not all proposals for additional distinctions appear to meet all criteria, and for some it is unclear whether they meet all criteria. Chapter 3 already stated that the segment-internal distinctions formulated in relation to the SOURCE OF COHERENCE primitive (presence of Subject of Consciousness [SoC], identity of the SoC, explicit mentioning of the SoC, and the propositional attitude of the segments) do not meet the relational criterion. The DIRECTNESS distinction within negative additive relations proposed by Pander Maat (1998) seems prone to the same problem discussed for negative additive versus negative causal relations above; although the DIRECTNESS distinction improves descriptive adequacy and some evidence for the distinction can be found in the Dutch connective system (see Pander Maat 1998:199), the cognitive plausibility of the DIRECTNESS distinction should be further investigated.

Chapter 3 also argued in favor of adding a new distinction to CCR, DISJUNCTION, to capture the difference between disjunctions (+*alternative*) and other types of relations (-*alternative*). It was argued that DISJUNCTION is a property of the relational surplus and that it applies to relations that hold between clauses or larger discourse segments. Evidence for the cognitive plausibility of DISJUNCTION was mainly found in linguistic systems; many languages have a connective that prototypically marks *disjunctions*. Since very little experimental work has focused on clausal disjunction (but see Staub & Clifton 2006), the processing evidence for the cognitive plausibility of the DISJUNCTION distinction is limited. The processing of *disjunction* relations thus appears to be a topic worth investigating in the future.

An issue with the status of DISJUNCTION that was very briefly mentioned in Chapter 3 is the similarity between *disjunctions* and *unless*-relations. While some annotation approaches analyze *unless*-relations as negative conditionals (Carlson & Marcu 2001, PDTB 3.0; Prasad, Webber, Lee, & Joshi *in prep* [see Rehbein, Scholman, & Demberg 2016 for the annotation scheme]), others group them with *disjunctions* (PDTB 2.0; PDTB Research Group 2007, Reese, Hunter, Asher, Denis, & Baldridge 2007). Both *disjunctions* and *unless*-relations contain alternative scenarios (see also Dancygier 1985) and *unless*-relations can often be expressed with *or* instead of *unless* and retain its overall meaning (*you either know it or you don't // unless you know it, you don't*). Since *unless*-relations were classified as negative conditional relations in CCR before the introduction of the DISJUNCTION distinction, Chapter 3 stated that this should continue to be the case. The exact differences and similarities between *disjunctions* and *unless*-relations, however, should receive more (theoretical) consideration to determine how *unless*-relations should be treated within the field of discourse in general, and CCR in specific.

An issue that was left unaddressed in Chapter 3 is whether two discourse segments can be related to each other by more than one coherence relation. In discourse annotation, coders are usually asked to pick a single relation label, but this is not always an easy task; a relation may for instance be ambiguous between two relation types or a coder might have the impression that two relations hold at the same time. For example, it might be difficult to choose between a positive objective additive relation with a synchronous temporal order or a negative additive objective relation when annotating a relation marked by English *while* (e.g., *While I'm cleaning the kitchen, you're watching TV*). In such cases, discourse annotation approaches can ask a coder to choose the most prominent interpretation (e.g., Sanders & Spooren 1999), or they can allow double-tagging of the relation (e.g., PDTB Research Group 2007). Recent empirical studies provide some evidence that multiple relations holding between two discourse segments is indeed a possibility (Rohde et al. 2016, Scholman & Demberg 2017). More research into the cognitive plausibility of two discourse segments being related to each other by more than one coherence relation cannot only help formulate clear guidelines for double-tagging when using CCR for discourse annotation, but also further our understanding of how language users construct a coherent mental representation of a discourse.

As this dissertation has shown, the primitives of the Cognitive approach to Coherence Relations are very much suited to use as a basis for discourse annotation. Depending on the research purpose, the taxonomy can easily be extended with additional features that improve the descriptive adequacy of the annotation scheme. For new distinctions to be adopted into CCR, however, they have to be cognitively plausible. By continuing to study the cognitive plausibility of existing primitives, combinations of primitive values, and potential new distinctions, CCR can be further developed into a model of discourse coherence that is both cognitively plausible and descriptively adequate (see also Scholman *in prep* for a more elaborate discussion on cognitive plausibility versus descriptive adequacy).

7.5 Conclusion

This dissertation focused on discourse segmentation and the linguistic marking of coherence relations. It provided new insights into how and why coherence relations are explicitly signaled and between which parts of a text people infer coherence relations. In addition, it contributed toward refining discourse segmentation and annotation guidelines, both of which are important methodological tools in the research of discourse coherence. The continued study of coherence relations and discourse structure is essential in furthering our understanding of how language users produce and comprehend discourse, a crucial component of human communication.

Appendix A Source of examples

All examples with an ep-number were taken from the Europarl corpus (Koehn 2005).

All examples without an ep-number that are not listed here and not explicitly cited as an example from another paper were created for the purpose of illustrating the linguistic phenomena discussed in this dissertation.

Chapter 1

- (1) Christensen, E (2015, June 22). *How to make no-cook, no-churn, 2-ingredient ice cream*. Retrieved from <https://www.thekitchn.com/how-to-make-no-churn-ice-cream-cooking-lessons-from-the-kitchn-220413>
- (2) Schmidt, K. (2017, July 28). *Here's how long ice cream can last in the freezer (and how to store it)*. Retrieved from <https://www.thekitchn.com/5-tips-for-storing-ice-cream-and-preventing-freezer-burn-246502>
- (3) Sugar, R. (2017, April 24). *In defense of keeping stuff you never use*. Retrieved from <https://www.thekitchn.com/do-you-need-an-ice-cream-maker-probably-243511>

Chapter 2

- (32) *Tripadvisor review* (2009, August 31). Retrieved from http://www.tripadvisor.co.uk/ShowUserReviews-g211878-d654059-r39310708-Cloisters_Bed_Break_fastKinsale_County_Cork.html

Chapter 3

- (1) Judkis, M. (2018, February 5). Doritos is developing lady-friendly chips because you should never hear a woman crunch. *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/food/wp/2018/02/05/doritos-is-developing-lady-friendly-chips-because-apparently-you-should-never-hear-a-woman-crunch/?utm_term=.f0f41dcb8328
- (2) Sedaris, D. (1997). *Naked*. New York, NY: Little, Brown and Company. p.276
- (3) Jackson, S. (1962). *We Have Always Lived in the Castle*. New York, NY: Viking Press. p.21
- (4) Tchou, W. (2017, December 11). Tokyo record bar's riff on the speakeasy. *The New Yorker*. Retrieved from <https://www.newyorker.com/magazine/2017/12/11/tokyo-record-bars-riff-on-the-speakeasy>
- (5) Crane, D., Kauffman, M., Astrof, J., Sikowitz, M., Chase, A., & Ungerleider, I. (Writers) & Lazarus, P. (Director). *The one with the dozen lasagnas*, S01E12

- [television series episode]. In Crane, D. & Kauffman, M. (Creators), *Friends*. USA: Warner Bros. Television & Bright/Kauffman/Crane Productions.
- (6) Foster, G. (Producer) & Ephron, N. (Director). *Sleepless in Seattle* [motion picture].
 - (7) Krauss, N. (2005). *The History of Love*. New York, NY: W.W. Norton & Company. p.77
 - (8) Anderson, W. (Producer & Director). (2012). *Moonrise Kingdom* [motion picture].
 - (9) 'Wee' Harry Potter fest cancelled because of popularity (2017, March 29). Retrieved from <http://www.bbc.co.uk/newsbeat/article/39431618/wee-harry-potter-fest-cancelled-because-of-popularity>
 - (10) 8 things you could knit/crochet for your wedding (2014, May 7). Retrieved from <http://ihatecleaning.com.au/8-things-you-could-knitcrochet-for-your-wedding/>
 - (11) S9E57 [television series episode]. In Roddam, F. (Creator), *Masterchef Australia*. Australia: FremantleMedia Australia, FremantleMedia, & Shine Australia.
 - (12) Ng, C. (2017). *Little fires everywhere*. London: Penguin Press. p.14
 - (13) Krauss, N. (2005). *The History of Love*. New York, NY: W.W. Norton & Company. p.104
 - (14) Penguin Point (2018). Retrieved from <http://www.vanaqua.org/experience/exhibits-and-galleries/penguin-point>
 - (15) Crane, D., Kauffman, M., & Curtis, M. (Writers) & Jensen, S. (Director). The one where Phoebe hates PBS, S05E04 [television series episode]. In Crane, D. & Kauffman, M. (Creators), *Friends*. USA: Warner Bros. Television & Bright/Kauffman/Crane Productions.
 - (16) Miller, C. (2018, January 9). *Jimmy Iovine denies he's leaving Apple, says he wants to make streaming 'more interesting.'* Retrieved from <https://9to5mac.com/2018/01/09/jimmy-iovine-not-leaving-apple/>
 - (17) Casciano, M. (2018, February). *7 cheapest beach vacations if you want your Instagram to sizzle*. Retrieved from <https://www.elitedaily.com/p/7-cheapest-beach-vacations-if-you-want-your-instagram-to-sizzle-8174172>
 - (18) Luker, G. (2014, August). *How to use silhouette software to cut your own graphics*. Retrieved from <https://www.theshabbycreekcottage.com/use-silhouette-software-cut-graphics.html>
 - (19) Last, T.S. (2018, February 15). *Mayoral election: 'The game has changed.'* Retrieved from <https://www.abqjournal.com/1134349/mayoral-election-the-game-has-changed-ex-while-knocking-on-doors-is-still-part-of-a-campaign-social-media-is-vital-today.html>
 - (20) Crane, D., Kauffman, M., & Abrams, D. (Writers) & Mancuso, G. (Director). The one where Ross can't flirt, S05E19 [television series episode]. In Crane, D.

- & Kauffman, M. (Creators), *Friends*. USA: Warner Bros. Television & Bright/Kauffman/Crane Productions.
- (21) *104 extra beds for rough sleepers during cold snap* (2018, February 26). Retrieved from <https://www.rte.ie/news/weather/2018/0226/943592-weather-snow/>
- (22) Lozada, C. (2015, June 4). From ‘Jill’ to ‘Mom’ – inside Jill Biden’s relationship with Beau and Hunter. *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/book-party/wp/2015/06/04/from-jill-to-mom-inside-jill-bidens-relationship-with-beau-and-hunter/?utm_term=.da2586370201
- (23) Guadiano, A.M. & Vlastelica, R. (2018, February 12). *Stock market notches daily gain, but posts largest weekly drop since early 2016*. Retrieved from <https://www.marketwatch.com/story/us-stock-futures-rise-as-dow-faces-worst-week-since-the-global-financial-crisis-2018-02-09>
- (24) Kirshner, R. (Writer) & Clancy, S. (Director), Hay bale maze, S7E18 [television series episode]. In Sherman-Palladino, A. (Creator), *Gilmore Girls*. USA: Dorothy Parker Drank Here Productions, Hofflund/Polone, Warner Bros. Television.
- (25) Hill, N. (2016). *The Nix*. London: Picador. p.337
- (26) Sedaris, D. (2013). *Let’s Explore Diabetes with Owls*. New York, NY: Little, Brown and Company. p. 21.
- (27) Grinsteiner, K. (2018, April 13). ‘You either know it, or you don’t.’ Retrieved from http://www.hibbingmn.com/news/local/you-either-know-it-or-you-dont/article_17178542-3ebe-11e8-a10e-5ba84094a972.html
- (28) Brant, K. (2018, February 28). Wayward waffles: Cornmeal and nontraditional additives set these bad boys apart from their all-flour relatives. *Arkansas Online*. Retrieved from <http://www.arkansasonline.com/news/2018/feb/28/wayward-waffles-20180228/>
- (29) *Philippines*. (2016). Retrieved from https://www.reddit.com/r/Philippines/comments/3k4wr7/alleged_fake_uber_car_roaming_around_the_metro/cuv20ju/
- (30) *Are you just feeling lazy or do you need a break?* (2008, December 15). Retrieved from <https://www.dumblittleman.com/are-you-just-feeling-lazy-or-do-you/>
- (31) *Petaluma, California* (2017). Retrieved from <https://www.trover.com/d/1EwTK-petaluma-california>
- (32) Borowitz, A. (2017, October 6). DeVos defends Trump: “Would a moron hire me?” *The New Yorker*. Retrieved from <https://www.newyorker.com/humor/borowitz-report/devos-defends-trump-would-a-moron-hire-me>
- (33) Burroughs, E.R. (2012). *Tarzan of the Apes*. Retrieved from <https://www.cs.cmu.edu/~rgs/tarz-10.html>. Chapter 10.

- (34) Andrews, R. (2017, November 29). *Scientist calls for glitter to be banned because it's awful for the environment*. Retrieved from <http://www.iflscience.com/environment/scientist-calls-glitter-banned-awful-environment/>
- (35) Casey, B. (2012, August 30). *I had a face tattoo for a week*. Retrieved from https://www.vice.com/en_us/article/ppqm7v/i-had-a-face-tattoo-for-a-week
- (37) Thomson, J.R. (2014, June 13). *13 things you should never put in the microwave*. Retrieved from https://www.huffingtonpost.com/2014/06/13/microwave-cooking-tips_n_5488231.html
- (38) Hanson, C. (2017). Why cooking with wine makes food taste better. Retrieved from <http://dish.allrecipes.com/cooking-wine-makes-food-taste-better/>
- (39) Korsh, A. & Cowan, J. (Writers) & Kumble, R (Director). I want you to want me, S03E02 [television series episode]. In Aaron Korsh (Creator), *Suits*. USA: Hypnotic, Universal Cable Productions, & Dutch Oven.

Chapter 6

- (1) Man who attacked jogger in Seattle park sentenced to prison (2018, March 24). *Seattle Times*. Retrieved from <https://www.seattletimes.com/seattle-news/man-who-attacked-jogger-in-seattle-park-sentenced-to-prison/>
- (2) Hill, N. (2016). *The Nix*. London: Picador. p.36
- (3) Addison, Bill (2018, March 7). *The state of dining in Texas*. Retrieved from <https://www.eater.com/2018/3/7/17080432/texas-dining-barbecue-mexican-crawfish>

Chapter 7

- (1) Claire (2018, April 5). Re: Are you becoming your parents? [blog post comment]. Retrieved from <https://cupofjo.com/2018/04/becoming-your-parents/#comments>
- (2) Nancy (2018, April 7). Re: Five words that changed everything [blog post comment]. Retrieved from <https://cupofjo.com/2018/04/five-words-that-changed-everything/#comments>
- (3) Sarah (2018, April 4). Re: Do or don't: Very scary books [blog post comment]. Retrieved from <https://cupofjo.com/2018/04/scariest-true-crime-thriller/#comments>
- (4) Velden, Dana (2018, April 9). *How to make an upside-down cake with almost any fruit*. Retrieved from <https://www.thekitchn.com/how-to-make-an-upside-down-cake-with-almost-any-fruit-cooking-lessons-from-the-kitchn-192070>
- (5) Michelle (2018, March 23). Re: Do or don't: Sharing food at restaurants [blog post comment]. Retrieved from <https://cupofjo.com/2018/03/i-dont-like-sharing-at-restaurants/#comments>

- (6) Aimee (2018, March 2). Re: A morning adventure in Brooklyn [blog post comment]. <https://cupofjo.com/2018/03/brooklyn-with-kids/#comments>
- (7) *Cassie Hayes: Man jailed for murdering travel agent* (2018, April 9). Retrieved from <http://www.bbc.com/news/uk-england-merseyside-43699004>
- (8) Abrahamson, E. (2018, February 11). Real lives: 'I was 30, single, with a successful career when I fell in love with a man who lived in a bush.' *Daily Mail Online*. Retrieved from <http://www.dailymail.co.uk/home/you/article-5330933/Real-lives-fell-love-homeless-man.html>
- (9) Strout, E. (2008). *Olive Kitteridge*. New York, NY: Random House. p. 124
- (10) Cheever, J. (1977). *Falconer*. New York, NY: Alfred A. Knopf. p.89
- (11) Sedaris, D. (2017). *Theft by Finding*. New York, NY: Little, Brown and Company. p.17
- (12) Sedaris, D. (2000). *Me Talk Pretty One Day*. New York, NY: Little, Brown and Company. p.182
- (13) Sedaris, D. (2000). *Me Talk Pretty One Day*. New York, NY: Little, Brown and Company. p.7-8

Appendix B Experimental items Experiments 1 and 2, Chapter 6

In preparation for Experiment 3 (self-paced reading), all items were constructed in pairs. The main verbs in each pair (1-2, 3-4, etc.) function as antonyms; the same RC is used in the causal condition for the one verb and in the concessive condition for the other verb. Each list contained only one instance of each main verb and an RC occurred only once in each list. Each prompt ended with either *because* or *even though* (Experiment 1) or with *because* (Experiment 2).

1. [neutr] Andrew admired the woman who came to visit him last week
[caus] Andrew admired the woman who had built a successful career in sales
[conc] Andrew admired the woman who had lost four jobs within a year
2. [neutr] Michael pitied the woman who came to visit him last week
[caus] Michael pitied the woman who had lost four jobs within a year
[conc] Michael pitied the woman who had built a successful career in sales
3. [neutr] Susan praised the guy who is always wearing a blue shirt
[caus] Susan praised the guy who made a lot of money for the company
[conc] Susan praised the guy who was 30 minutes late for the meeting
4. [neutr] Charlotte fired the guy who is always wearing a blue shirt
[caus] Charlotte fired the guy who was 30 minutes late for the meeting
[conc] Charlotte fired the guy who made a lot of money for the company
5. [neutr] James congratulated the actress who was wearing yellow shoes
[caus] James congratulated the actress who had won an Oscar
[conc] James congratulated the actress who had quit halfway through the marathon
6. [neutr] Henry criticized the actress who was wearing yellow shoes
[caus] Henry criticized the actress who had quit halfway through the marathon
[conc] Henry criticized the actress who had won an Oscar
7. [neutr] We thanked the neighbor who stopped by on Tuesday night
[caus] We thanked the neighbor who brought over a fruit basket
[conc] We thanked the neighbor who dropped our newly inherited vase
8. [neutr] We sued the neighbor who stopped by on Tuesday night
[caus] We sued the neighbor who dropped our newly inherited vase
[conc] We sued the neighbor who brought over a fruit basket
9. [neutr] Mrs. Miller valued the gardener who took special care of her tulips
[caus] Mrs. Miller valued the gardener who took special care of her tulips
[conc] Mrs. Miller valued the gardener who never took off his muddy shoes
10. [neutr] Mrs. Thompson loathed the gardener who came by every Wednesday
[caus] Mrs. Thompson loathed the gardener who never took off his muddy shoes
[conc] Mrs. Thompson loathed the gardener who took special care of her tulips
11. [neutr] I complimented the child who just moved here from Scandinavia
[caus] I complimented the child who had gotten a perfect test score
[conc] I complimented the child who had thrown a pair of scissors
12. [neutr] I scolded the child who just moved here from Scandinavia
[caus] I scolded the child who had thrown a pair of scissors
[conc] I scolded the child who had gotten a perfect test score

13. [neutr] Emily comforted the man who was sitting in the grey Volvo
 [caus] Emily comforted the man who was crying on the bus
 [conc] Emily comforted the man who was making a mess in the waiting room
14. [neutr] Mia yelled at the man who was sitting in the grey Volvo
 [caus] Mia yelled at the man who was making a mess in the waiting room
 [conc] Mia yelled at the man who was crying on the bus
15. [neutr] Emma trusted the doctor who would be setting her leg
 [caus] Emma trusted the doctor who had received a prestigious award
 [conc] Emma trusted the doctor who had messed up the procedure last time
16. [neutr] Natalie distrusted the doctor who would be setting her leg
 [caus] Natalie distrusted the doctor who had messed up the procedure last time
 [conc] Natalie distrusted the doctor who had received a prestigious award
17. [neutr] I noticed the girl who was standing in the doorway
 [caus] I noticed the girl who was wearing neon pink glitter leggings
 [conc] I noticed the girl who was hiding in the corner
18. [neutr] I overlooked the girl who was standing in the doorway
 [caus] I overlooked the girl who was hiding in the corner
 [conc] I overlooked the girl who was wearing neon pink glitter leggings
19. [neutr] Aaron respected the aunt who lived on the other side of the country
 [caus] Aaron respected the aunt who worked for Doctors without Borders
 [conc] Aaron respected the aunt who did not recycle
20. [neutr] Oliver condemned the aunt who lived on the other side of the country
 [caus] Oliver condemned the aunt who did not recycle
 [conc] Oliver condemned the aunt who worked for Doctors without Borders
21. [neutr] Mr. Johnson adored the ballerina who had joined the group last year
 [caus] Mr. Johnson adored the ballerina who practiced hard every day
 [conc] Mr. Johnson adored the ballerina who never listened to the instructions
22. [neutr] Mr. Smith despised the ballerina who had joined the group last year
 [caus] Mr. Smith despised the ballerina who never listened to the instructions
 [conc] Mr. Smith despised the ballerina who practiced hard every day
23. [neutr] Ms. Walker rewarded the boy who was sitting by the window
 [caus] Ms. Walker rewarded the boy who swept the classroom floor
 [conc] Ms. Walker rewarded the boy who had stolen a pillow case
24. [neutr] Ms. Roberts punished the boy who was sitting by the window
 [caus] Ms. Roberts punished the boy who had stolen a pillow case
 [conc] Ms. Roberts punished the boy who swept the classroom floor
25. [neutr] Suzie loved the uncle who lived across the street
 [caus] Suzie loved the uncle who often bought her flowers
 [conc] Suzie loved the uncle who did not show up for her 30th birthday party
26. [neutr] Lauren hated the uncle who lived across the street
 [caus] Lauren hated the uncle who did not show up for her 30th birthday party
 [conc] Lauren hated the uncle who often bought her flowers
27. [neutr] Paul envied the stewardess who was walking down the aisle
 [caus] Paul envied the stewardess who got to stay in Paris for the weekend
 [conc] Paul envied the stewardess who crashed the drink cart into the wall
28. [neutr] Geoff ridiculed the stewardess who was walking down the aisle
 [caus] Geoff ridiculed the stewardess who crashed the drink cart into the wall
 [conc] Geoff ridiculed the stewardess who got to stay in Paris for the weekend
29. [neutr] Caroline applauded the congressman who was voted in seven years ago
 [caus] Caroline applauded the congressman who passed the bipartisan bill
 [conc] Caroline applauded the congressman who lost the most recent election

30. [neutr] Alice reassured the congressman who was voted in seven years ago
[caus] Alice reassured the congressman who lost the most recent election
[conc] Alice reassured the congressman who passed the bipartisan bill

Appendix C Experimental items Experiment 3, Chapter 6

All items were constructed in pairs. The main verbs in the causal and concessive conditions in each pair (1-2, 3-4, etc.) function as antonyms; the same RC is used in the causal condition for the one verb and in the concessive condition for the other verb. Main verbs and RCs occurred only once in each list.

1.

Intro	Andrew looked over the crowd that had assembled in the company lounge.
<i>neutral RC</i>	He talked to the woman who had built a successful career in sales.
<i>causal RC</i>	He admired the woman who had built a successful career in sales.
<i>concessive RC</i>	He pitied the woman who had built a successful career in sales.
Wrap-up	She arrived at the conference room just in time for her next meeting.
Verification statement	<i>The woman has a successful career.</i>

2.

Intro	Last week, George attended a huge career event, where he met a lot of new people.
<i>neutral RC</i>	He talked to the woman who had lost four jobs within a year.
<i>causal RC</i>	He pitied the woman who had lost four jobs within a year.
<i>concessive RC</i>	He admired the woman who had lost four jobs within a year.
Wrap-up	He registered for three follow-up courses in the upcoming weeks.
Verification statement	<i>The woman lost four jobs within a year.</i>

3.

Intro	Jenny walked through the hallway to check on the daily goings-on around the office.
<i>neutral RC</i>	She joked with the guy who made a lot of money for the company.
<i>causal RC</i>	She praised the guy who made a lot of money for the company.
<i>concessive RC</i>	She fired the guy who made a lot of money for the company.
Wrap-up	She arrived at the conference room just in time for her next meeting.
Verification statement	<i>The guy made a lot of money for the company.</i>

4.

Intro	Charlotte was almost ready to close the weekly company assembly.
<i>neutral RC</i>	She joked with the guy who had been 30 minutes late to the meeting.
<i>causal RC</i>	She fired the guy who had been 30 minutes late to the meeting.
<i>concessive RC</i>	She praised the guy who had been 30 minutes late to the meeting.
Wrap-up	She then announced that it was time for drinks.
Verification statement	<i>The guy was late for the meeting.</i>

5.

Intro	James walked into the busy lecture hall.
<i>neutral RC</i>	He waved at the girl who had won the writing contest.
<i>causal RC</i>	He congratulated the girl who had won the writing contest.
<i>concessive RC</i>	He criticized the girl who had won the writing contest.
Wrap-up	He picked a seat near the front and sat down.
Verification statement	<i>The girl won the writing contest.</i>

6.

Intro	Yesterday, Kyle went to the city center to watch the marathon with his friends.
<i>neutral RC</i>	He waved at the girl who had quit half-way through the race.
<i>causal RC</i>	He criticized the girl who had quit half-way through the race.
<i>concessive RC</i>	He congratulated the girl who had quit half-way through the race.
Wrap-up	After the event, he had dinner at an Italian restaurant.
Verification statement	<i>The girl quit halfway through a marathon.</i>

7.

Intro	We entered our apartment building and headed up the stairs.
<i>neutral RC</i>	We greeted the neighbor who had brought over a fruit basket.
<i>causal RC</i>	We thanked the neighbor who had brought over a fruit basket.
<i>concessive RC</i>	We ignored the neighbor who had brought over a fruit basket.
Wrap-up	When we reached our apartment, we discovered we left our keys in the front door.
Verification statement	<i>Our neighbor brought us a fruit basket.</i>

8.

Intro	As we left the house, we told our son to join us at the neighborhood picnic soon.
<i>neutral RC</i>	We greeted the neighbor who had dropped our newly inherited vase.
<i>causal RC</i>	We ignored the neighbor who had dropped our newly inherited vase.
<i>concessive RC</i>	We thanked the neighbor who had dropped our newly inherited vase.
Wrap-up	We put our homemade egg salad on one of the tables.
Verification statement	<i>Our neighbor dropped our vase.</i>

9.

Intro	Having just finished marking the exams, I walked into the classroom.
<i>neutral RC</i>	I saw the child who had gotten a perfect test score.
<i>causal RC</i>	I complimented the child who had gotten a perfect test score.
<i>concessive RC</i>	I scolded the child who had gotten a perfect test score.
Wrap-up	I placed the stack of exams on my desk.
Verification statement	<i>The child got a perfect score.</i>

10.

Intro	Carrying a huge bowl of potato chips, I entered the room where the birthday party was taking place.
<i>neutral RC</i>	I saw the child who had thrown a pair of scissors.
<i>causal RC</i>	I scolded the child who had thrown a pair of scissors.
<i>concessive RC</i>	I complimented the child who had thrown a pair of scissors.
Wrap-up	I put the bowl down on one of the empty chairs.
Verification statement	<i>The child threw a pair of scissors.</i>

11.

Intro	During her visit to Morningside Hospital, Emma got to meet almost the entire hospital staff.
<i>neutral RC</i>	She knew the doctor who had received a prestigious award.
<i>causal RC</i>	She trusted the doctor who had received a prestigious award.
<i>concessive RC</i>	She distrusted the doctor who had received a prestigious award.
Wrap-up	She was visiting one more hospital before deciding where she would do her residency.
Verification statement	<i>The doctor has received an award.</i>

12.

Intro	Natalie had to go to the hospital soon to have her cast removed.
<i>neutral RC</i>	She knew the doctor who had messed up the procedure last time.
<i>causal RC</i>	She distrusted the doctor who had messed up the procedure last time.
<i>concessive RC</i>	She trusted the doctor who had messed up the procedure last time.
Wrap-up	She thought of everything she would be able to do again when she had her left arm back.
Verification statement	<i>Natalie's doctor had messed up her procedure.</i>

13.

Intro	Aaron kept in touch with most of his mother's seven sisters.
<i>neutral RC</i>	He carpooled with the aunt who worked for Doctors without Borders.
<i>causal RC</i>	He respected the aunt who worked for Doctors without Borders.
<i>concessive RC</i>	He condemned the aunt who worked for Doctors without Borders.
Wrap-up	He recently heard that two of the other aunts had moved to Canada.
Verification statement	<i>Aaron's aunt works for Doctors without Borders.</i>

14.

Intro	Oliver had to attend his family reunion every single year.
<i>neutral RC</i>	He carpooled with the aunt who had punched her care-taker on several occasions.
<i>causal RC</i>	He condemned the aunt who had punched her care-taker on several occasions.
<i>concessive RC</i>	He respected the aunt who had punched her care-taker on several occasions.
Wrap-up	He stopped going when he moved abroad to live with his girlfriend.
Verification statement	<i>Oliver's aunt punched her care-taker.</i>

15.

Intro	Ms. Roberts entered the school yard.
<i>neutral RC</i>	She recognized the boy who had swept the classroom floor.
<i>causal RC</i>	She rewarded the boy who had swept the classroom floor.
<i>concessive RC</i>	She punished the boy who had swept the classroom floor.
Wrap-up	She headed into the building to prepare for her first class.
Verification statement	<i>The boy swept the classroom floor.</i>

16.

Intro	Ms. Walker was called into one of her colleague's classrooms.
<i>neutral RC</i>	She recognized the boy who had stolen a pencil case.
<i>causal RC</i>	She punished the boy who had stolen a pencil case.
<i>concessive RC</i>	She rewarded the boy who had stolen a pencil case.
Wrap-up	She returned to her own class just in time to prevent a fight between two girls in the front row.
Verification statement	<i>The boy stole a pencil case.</i>

17.

Intro	Ginny opened the door to let in three of her father's brothers.
<i>neutral RC</i>	She resembled the uncle who always bought amazing presents.
<i>causal RC</i>	She loved the uncle who always bought amazing presents.
<i>concessive RC</i>	She hated the uncle who always bought amazing presents.
Wrap-up	She quickly ran to the kitchen to take the kettle off the stove.
Verification statement	<i>Ginny's uncle buys great presents.</i>

18.

Intro	At dinner, Suzie told her best friend about her family members.
<i>neutral RC</i>	She resembled the uncle who forgot about her 30th birthday party.
<i>causal RC</i>	She hated the uncle who forgot about her 30th birthday party.
<i>concessive RC</i>	She loved the uncle who forgot about her 30th birthday party.
Wrap-up	She had not seen that uncle in almost a year.
Verification statement	<i>Suzie's uncle forgot her birthday party.</i>

19.

Intro	Caroline attended an election rally in her home state.
<i>neutral RC</i>	She ran into the congressman who passed the bipartisan bill.
<i>causal RC</i>	She applauded the congressman who passed the bipartisan bill.
<i>concessive RC</i>	She reassured the congressman who passed the bipartisan bill.
Wrap-up	She wrestled her way through the crowd and stood right in front of the stage for most of the event.
Verification statement	<i>The congressman passed a bill.</i>

20.

Intro	On her trip to Washington D.C., Tara spent a lot of time around politicians.
<i>neutral RC</i>	She ran into the congressman who lost the most recent election.
<i>causal RC</i>	She reassured the congressman who lost the most recent election.
<i>concessive RC</i>	She applauded the congressman who lost the most recent election.
Wrap-up	On the last day, she saw the Vice-President walking out of a restaurant.
Verification statement	<i>The congressman lost the election.</i>

21.

Intro	Jake looked out of the airplane window and noticed they were already flying over the Atlantic Ocean.
<i>neutral RC</i>	He talked to the stewardess who got to stay in Paris for the weekend.
<i>causal RC</i>	He envied the stewardess who got to stay in Paris for the weekend.
<i>concessive RC</i>	He ridiculed the stewardess who got to stay in Paris for the weekend.
Wrap-up	He ordered a drink and selected a movie from the entertainment system's menu.
Verification statement	<i>The stewardess spent the weekend in Paris.</i>

22.

Intro	After dropping off his suitcase and passing through security, Geoff headed for the gate.
<i>neutral RC</i>	He talked to the stewardess who crashed the drink cart into the wall.
<i>causal RC</i>	He ridiculed the stewardess who crashed the drink cart into the wall.
<i>concessive RC</i>	He envied the stewardess who crashed the drink cart into the wall.
Wrap-up	Boarding began just five minutes after the announced boarding time.
Verification statement	<i>The stewardess crashed the drink cart into the wall.</i>

23.

Intro	At the annual charity event, Mrs. Miller talked freely about her relationship with her staff.
<i>neutral RC</i>	She gossiped with the gardener who took special care of her tulips.
<i>causal RC</i>	She valued the gardener who took special care of her tulips.
<i>concessive RC</i>	She loathed the gardener who took special care of her tulips.
Wrap-up	She got more than a little tipsy that night.
Verification statement	<i>The gardener takes good care of the tulips.</i>

24.

Intro	Lady Thompson was hosting an afternoon of bridge while her landscaping crew worked in the garden.
<i>neutral RC</i>	She gossiped with the gardener who always traipsed his muddy shoes through the house.
<i>causal RC</i>	She loathed the gardener who always traipsed his muddy shoes through the house.
<i>concessive RC</i>	She valued the gardener who always traipsed his muddy shoes through the house.
Wrap-up	The get-together was briefly in danger when it appeared that they had run out of brandy.
Verification statement	<i>The gardener never takes off his muddy shoes</i>

Appendix D Experimental items Experiment 4, Chapter 6

1.

Intro	We were recently involved in a law suit.
<i>neutral RC</i>	We were suing the neighbours from upstairs who moved in last year
+ <i>because</i>	because suddenly they decided it was okay to smoke in the building.
<i>neutral RC</i>	We were suing the neighbours from upstairs who moved in last year and so
+ <i>and so</i>	suddenly they decided we were unfriendly.
<i>causal RC</i>	We were suing the neighbours from upstairs who flooded our flat last year
+ <i>because</i>	because suddenly they decided it was okay to smoke in the building.
<i>causal RC</i>	We were suing the neighbours from upstairs who flooded our flat last year
+ <i>and so</i>	and so suddenly they decided we were unfriendly.
Wrap-up	The hearing took place last Monday.

2.

Intro	The company meeting ended with a few announcements.
<i>neutral RC</i>	Susan praised the accountant who is always wearing a blue shirt because
+ <i>because</i>	finally he realized that he must show up for meetings on time.
<i>neutral RC</i>	Susan praised the accountant who is always wearing a blue shirt and so
+ <i>and so</i>	finally he realized that people do actually notice his hard work.
<i>causal RC</i>	Susan praised the accountant who is always working overtime because
+ <i>because</i>	finally he realized that he must show up for meetings on time.
<i>causal RC</i>	Susan praised the accountant who is always working overtime and so
+ <i>and so</i>	finally he realized that people do actually notice his hard work.
Wrap-up	Everyone applauded politely.

3.

Intro	Let me fill you in on the latest company gossip.
<i>neutral RC</i> +	Diane fired the guy from the London office who was here last month
<i>because</i>	because astoundingly he hired a stripper for the Christmas party.
<i>neutral RC</i> +	Diane fired the guy from the London office who was here last month and
<i>and so</i>	so astoundingly he hired a lawyer to sue the company.
<i>causal RC</i> +	Diane fired the guy from the London office who was embezzling money
<i>because</i>	because astoundingly he hired a stripper for the Christmas party.
<i>causal RC</i> +	Diane fired the guy from the London office who was embezzling money
<i>and so</i>	and so astoundingly he hired a lawyer to sue the company.
Wrap-up	Also, Harold in accounting has received a promotion.
<i>Verification statement</i>	<i>Harold received a promotion.</i>

4.

Intro	Mr. Fitzgerald was a real grouch.
<i>neutral RC</i> + <i>because</i>	He criticized the girl from next door who was on the swimming team because sometimes she would ignore the stop sign at the end of their street.
<i>neutral RC</i> + <i>and so</i>	He criticized the girl from next door who was on the swim team and so sometimes she would ignore him when they ran into each other on the street.
<i>causal RC</i> + <i>because</i>	He criticized the girl from next door who often smoked in her bedroom because sometimes she would ignore the stop sign at the end of their street.
<i>causal RC</i> + <i>and so</i>	He criticized the girl from next door who often smoked in her bedroom and so sometimes she would ignore him when they ran into each other on the street.
Wrap-up	He also made mean comments whenever the local kids were playing outside.

5.

Intro	Paul loved the clothing shop in his home town.
<i>neutral RC</i> + <i>because</i>	He really valued the lady at the shop who sat behind the counter because often she would sense if he was in need of a compliment on his outfit.
<i>neutral RC</i> + <i>and so</i>	He really valued the lady at the shop who sat behind the counter and so often she would sense that he really appreciated her help.
<i>causal RC</i> + <i>because</i>	He really valued the lady at the shop who helped him find his favourite suit because often she would sense if he was in need of a compliment on his outfit.
<i>causal RC</i> + <i>and so</i>	He really valued the lady at the shop who helped him find his favourite suit and so often she would sense that he really appreciated her help.
Wrap-up	Unfortunately, the store closed in June.

6.

Intro	Yesterday, Natalie had to go to hospital.
<i>neutral RC</i> + <i>because</i>	She distrusted the doctor who would be setting her leg because supposedly he received the worst patient reviews in the district.
<i>neutral RC</i> + <i>and so</i>	She distrusted the doctor who would be setting her leg and so supposedly he received several questions about the procedure he was following.
<i>causal RC</i> + <i>because</i>	She distrusted the doctor who had messed up the procedure last time because supposedly he received the worst patient reviews in the district.
<i>causal RC</i> + <i>and so</i>	She distrusted the doctor who had messed up the procedure last time and so supposedly he received several questions about the procedure he was following.
Wrap-up	She was allowed to go home just after lunch.
<i>Verification statement</i>	<i>Natalie had to go to hospital.</i>

7.

Intro	Roy was walking to the farmers' market on a sunny morning.
<i>neutral RC</i> + <i>because</i>	He congratulated the actress who lived across the street because again she was on the short-list for an award.
<i>neutral RC</i> + <i>and so</i>	He congratulated the actress who lived across the street and so again she was on the receiving end of a compliment.
<i>causal RC</i> + <i>because</i>	He congratulated the actress who had finished the marathon because again she was on the short-list for an award.
<i>causal RC</i> + <i>and so</i>	He congratulated the actress who had finished the marathon and so again she was on the receiving end of a compliment.
Wrap-up	He also greeted the local dog walker with six dogs in tow.
<i>Verification statement</i>	<i>Roy greeted the dog walker.</i>

8.

Intro	Today is the last day of school before Christmas.
<i>neutral RC</i> + <i>because</i>	I complimented the child who has recently moved here from Spain because obviously she has been having a hard time adjusting to our customs.
<i>neutral RC</i> + <i>and so</i>	I complimented the child who has recently moved here from Spain and so obviously she has been smiling from ear to ear for the last few minutes.
<i>causal RC</i> + <i>because</i>	I complimented the child who had gotten a perfect test score because obviously she has been having a hard time with her parents' divorce.
<i>causal RC</i> + <i>and so</i>	I complimented the child who had gotten a perfect test score and so obviously she has been smiling from ear to ear for the last few minutes.
Wrap-up	This afternoon, we will all sing carols and eat gingerbread cookies.

9.

Intro	Emily is having a tough time teaching Year Three this year.
<i>neutral RC</i> + <i>because</i>	Yesterday, she scolded the boy who sits directly in front of her because shockingly he threatened to pee on her desk.
<i>neutral RC</i> + <i>and so</i>	Yesterday, she scolded the boy who sits directly in front of her and so shockingly he threatened to pee on her desk.
<i>causal RC</i> + <i>because</i>	Yesterday, she scolded the boy who had thrown a pair of scissors because shockingly he threatened to pee on her desk.
<i>causal RC</i> + <i>and so</i>	Yesterday, she scolded the boy who had thrown a pair of scissors and so shockingly he threatened to pee on her desk.
Wrap-up	Summer break cannot come soon enough for Emily.

10.

Intro	Oliver was caught up in some family drama.
<i>neutral RC</i> + <i>because</i>	He had publicly condemned the wealthy aunt who lived in Scotland because reportedly she was in the possession of Nazi gold.
<i>neutral RC</i> + <i>and so</i>	He had publicly condemned the wealthy aunt who lived in Scotland and so reportedly she was in the process of cutting him out of her will.
<i>causal RC</i> + <i>because</i>	He had publicly condemned the wealthy aunt who had verbally abused her butler because reportedly she was in the possession of Nazi gold.
<i>causal RC</i> + <i>and so</i>	He had publicly condemned the wealthy aunt who verbally abused her butler and so reportedly she was in the process of cutting him out of her will.
Wrap-up	The upcoming family reunion was going to be an interesting event.

11.

Intro	Sally attended an election debate.
<i>neutral RC</i> + <i>because</i>	She enthusiastically applauded the politician who presented during the final hour because in the end he had the most inspiring message of the day.
<i>neutral RC</i> + <i>and so</i>	She enthusiastically applauded the politician who presented during the final hour of and so in the end he had the idea to invite her on stage.
<i>causal RC</i> + <i>because</i>	She enthusiastically applauded the politician who received a humanitarian award because in the end he had the most inspiring message of the day.
<i>causal RC</i> + <i>and so</i>	She enthusiastically applauded the politician who received a humanitarian award and so in the end he had the idea to invite her on stage.
Wrap-up	The event ended with a short performance by a local band.

12.

Intro	Geoff was on a red-eye flight to New York.
<i>neutral RC</i> + <i>because</i>	He ridiculed the stewardess who was walking down the aisle because clearly she refused to acknowledge that she needed a dress in a much larger size.
<i>neutral RC</i> + <i>and so</i>	He ridiculed the stewardess who was walking down the aisle and so clearly she refused to provide him with any kind of service for the rest of the flight.
<i>causal RC</i> + <i>because</i>	He ridiculed the stewardess who crashed the drink cart into one of the seats because clearly she refused to acknowledge that she needed a dress in a much larger size.
<i>causal RC</i> + <i>and so</i>	He ridiculed the stewardess who crashed the drink cart into one of the seats and so clearly she refused to provide him with any kind of service for the rest of the flight.
Wrap-up	The other passengers thought Geoff was incredibly rude.
<i>Verification statement</i>	<i>Geoff took a red-eye flight to New Mexico.</i>

13.

Intro	Ginny opened the door for three of her mother's brothers.
<i>neutral RC</i> + <i>because</i>	She hated the uncle who worked at her mother's company because for a long time he had been trying to claim the most valuable family heirloom.
<i>neutral RC</i> + <i>and so</i>	She hated the uncle who worked at her mother's company and so for a long time he had been the focal point of her withering stares.
<i>causal RC</i> + <i>because</i>	She hated the uncle who had run over her dog because for a long time he had been trying to claim the most valuable family heirloom.
<i>causal RC</i> + <i>and so</i>	She hated the uncle who had run over her dog and so for a long time he had been the focal point of her withering stares.
Wrap-up	She tried to keep her composure while inviting everyone into the living room.

14.

Intro	My son's eighth birthday party was a complete disaster.
<i>neutral RC</i> + <i>because</i>	Early on, I punished the girl who was dropped off by her grandpa because viciously she tried to steal one of the presents.
<i>neutral RC</i> + <i>and so</i>	Early on, I punished the girl who was dropped off by her grandpa and so viciously she tried to lock herself and another girl in one of the bedrooms.
<i>causal RC</i> + <i>because</i>	Early on, I punished the girl who knocked over the punch bowl because viciously she tried to steal one of the presents.
<i>causal RC</i> + <i>and so</i>	Early on, I punished the girl who knocked over the punch bowl and so viciously she tried to lock herself and another girl in one of the bedrooms.
Wrap-up	Later, it turned out the bowling alley had no record of our reservation.

15.

Intro	We try our best to maintain a good relationship with our community.
<i>neutral RC</i> + <i>because</i>	We repeatedly thanked the man who spoke at the Council meeting because yesterday he told us about an amazing investment opportunity.
<i>neutral RC</i> + <i>and so</i>	We repeatedly thanked the man who spoke at the Council meeting and so yesterday he told us that he appreciated being acknowledged.
<i>causal RC</i> + <i>because</i>	We repeatedly thanked the man who helped clean our gutters because yesterday he told us about an amazing investment opportunity.
<i>causal RC</i> + <i>and so</i>	We repeatedly thanked the man who helped clean our gutters and so yesterday he told us that he appreciated being acknowledged.
Wrap-up	Being nice to other people can be really rewarding.

202 Appendix D

16.

Intro	Andy met a lot of new people at the business fair last month
<i>neutral RC</i> + <i>because</i>	He unabashedly admired the woman who manned the stall in the corner because evidently she had been giving out great advice to everyone at the fair.
<i>neutral RC</i> + <i>and so</i>	He unabashedly admired the woman who manned the stall in the corner and so evidently she had been a bit embarrassed by the flattery.
<i>causal RC</i> + <i>because</i>	He unabashedly admired the woman who was very successful in her trade because evidently she had been giving out great advice to everyone at the fair.
<i>causal RC</i> + <i>and so</i>	He unabashedly admired the woman who was very successful in her trade and so evidently she had been a bit embarrassed by the flattery.
Wrap-up	He tried to tone down his enthusiasm for the remainder of the day.

17.

Intro	This morning, Pauline caught the bus to work.
<i>neutral RC</i> + <i>because</i>	She comforted the boy who was sitting next to her because undoubtedly he was very upset about something.
<i>neutral RC</i> + <i>and so</i>	She comforted the boy who was sitting next to her and so undoubtedly he was very appreciative of her caretaking nature.
<i>causal RC</i> + <i>because</i>	She comforted the boy who had just lost his mother because undoubtedly he was very upset about it.
<i>causal RC</i> + <i>and so</i>	She comforted the boy who had just lost his mother and so undoubtedly he was very appreciative of her caretaking nature.
Wrap-up	She ended up getting off three stops past her destination.
<i>Verification statement</i>	<i>Pauline got off at the correct bus stop.</i>

18.

Intro	Bob was conducting interviews for the shop assistant vacancy.
<i>neutral RC</i> + <i>because</i>	He clearly pitied the woman who was the first to come in because the whole time she had been nervously chewing her hair.
<i>neutral RC</i> + <i>and so</i>	He clearly pitied the woman who was the first to come in and so the whole time she had been extremely uncomfortable.
<i>causal RC</i> + <i>because</i>	He clearly pitied the woman who had a terrible resume because the whole time she had been nervously chewing her hair.
<i>causal RC</i> + <i>and so</i>	He clearly pitied the woman who had a terrible resume and so the whole time she had been extremely uncomfortable.
Wrap-up	He would certainly pick another candidate.

19.

Intro	Chrissy got into a bit of a fight at a party.
<i>neutral RC</i> + <i>because</i>	She laughed at the guy who was standing by the drinks because for a while he had entertained middle-aged rich women for money.
<i>neutral RC</i> + <i>and so</i>	She laughed at the guy who was standing by the drinks and so for a while he had entertained the idea of throwing her in the pool.
<i>causal RC</i> + <i>because</i>	She laughed at the guy who had slipped on a puddle of beer because for a while he had entertained middle-aged rich women for money.
<i>causal RC</i> + <i>and so</i>	She laughed at the guy who had slipped on a puddle of beer and so for a while he had entertained the idea of throwing her in the pool.
Wrap-up	She tried to apologize by getting him a slice of pizza.

20.

Intro	Mrs. Thompson had many family members she saw regularly.
<i>neutral RC</i> + <i>because</i>	She often worried about the nephew who lived in the apartment below because frequently he would try to hide empty wine bottles when she came over unexpectedly.
<i>neutral RC</i> + <i>and so</i>	She often worried about the nephew who lived in the apartment below and so frequently he would try to reassure her that he was doing fine.
<i>causal RC</i> + <i>because</i>	She often worried about the nephew who lived in the apartment below her because frequently he would try to hide empty wine bottles when she came over unexpectedly.
<i>causal RC</i> + <i>and so</i>	She often worried about the nephew who had been in rehab several times and so frequently he would try to reassure her that he was doing fine.
Wrap-up	She was also concerned about her sick sister.

21.

Intro	Prof. Roberts was in the middle of a lecture on global politics.
<i>neutral RC</i> + <i>because</i>	He corrected the girl who sat in the front row because curiously she started to claim that World War II never actually happened.
<i>neutral RC</i> + <i>and so</i>	He corrected the girl who sat in the front row and so curiously she started to argue with him about the geography of Europe.
<i>causal RC</i> + <i>because</i>	He corrected the girl who thought Paris was a country because curiously she started to claim that World War II never actually happened.
<i>causal RC</i> + <i>and so</i>	He corrected the girl who thought Paris was a country and so curiously she started to argue with him about the geography of Europe.
Wrap-up	He told her they would continue the discussion after class.

22.

Intro	After work, Mia went to a pub with some colleagues.
<i>neutral RC</i> + <i>because</i>	She told off the businessman who was having a drink at the bar because obnoxiously he voiced his sexist opinions about women in business.
<i>neutral RC</i> + <i>and so</i>	She told off the businessman who was having a drink at the bar and so obnoxiously he voiced his objection to her hostile attitude.
<i>causal RC</i> + <i>because</i>	She told off the businessman who spilled a drink on her dress because obnoxiously he voiced his sexist opinions about women in business.
<i>causal RC</i> + <i>and so</i>	She told off the businessman who spilled a drink on her dress and so obnoxiously he voiced his objection to her hostile attitude.
Wrap-up	After a while, she left to catch the last train home.

23.

Intro	Kim watched a few kids while their parents were at a school meeting.
<i>neutral RC</i> + <i>because</i>	She rewarded the boy who has a large mole on his cheek because for half an hour he sat quietly drawing cartoons.
<i>neutral RC</i> + <i>and so</i>	She rewarded the boy who has a large mole on his cheek and so for half an hour he sat quietly enjoying the praise.
<i>causal RC</i> + <i>because</i>	She rewarded the boy who had helped set up the crafts table because for half an hour he sat quietly drawing cartoons.
<i>causal RC</i> + <i>and so</i>	She rewarded the boy who had helped set up the crafts table and so for half an hour he sat quietly enjoying the praise.
Wrap-up	Some of the other kids were chasing after a squirrel.

24.

Intro	Grace was in her final year of high school.
<i>neutral RC</i> + <i>because</i>	She completely idolized the guy who sat in front of her during Art because regularly he would take his grandma to her yoga class.
<i>neutral RC</i> + <i>and so</i>	She completely idolized the guy who sat in front of her during Art and so regularly he would take advantage of her.
<i>causal RC</i> + <i>because</i>	She completely idolized the guy who was the cutest member of the hockey team because regularly he would take his grandma to her yoga class.
<i>causal RC</i> + <i>and so</i>	She completely idolized the guy who was the cutest member of the hockey team and so regularly he would take advantage of her.
Wrap-up	Her only real friend was Glenn, the school band's triangle player.
<i>Verification statement</i>	<i>Grace is in her last year of high school.</i>

25.

Intro	Mr. Brown was teaching his weekly sculpture class.
<i>neutral RC</i> + <i>because</i>	He chastised the lady who was making a modernist cube because allegedly she stole the idea for her project from another classmate.
<i>neutral RC</i> + <i>and so</i>	He chastised the lady who was making a modernist cube and so allegedly she stole the most expensive sculpting tool in retaliation.
<i>causal RC</i> + <i>because</i>	He chastised the lady who had dropped her clay for the seventh time because allegedly she stole the idea for her project from another classmate.
<i>causal RC</i> + <i>and so</i>	He chastised the lady who had dropped her clay for the seventh time and so allegedly she stole the most expensive sculpting tool in retaliation.
Wrap-up	Class ended at 4PM sharp.

26.

Intro	Queen Wendelyn was addressing her Court.
<i>neutral RC</i> + <i>because</i>	She banished the knight who guarded the castle's main entrance because astonishingly he tried to poison her the day before.
<i>neutral RC</i> + <i>and so</i>	She banished the knight who guarded the castle's main entrance and so astonishingly he tried to draw his sword and attack her.
<i>causal RC</i> + <i>because</i>	She banished the knight who had recently let three prisoners escape because astonishingly he tried to poison her the day before.
<i>causal RC</i> + <i>and so</i>	She banished the knight who had recently let three prisoners escape and so astonishingly he tried to draw his sword and attack her.
Wrap-up	Later, she also reprimanded the court jester.

27.

Intro	Mr. Evans was having lunch at his favourite pub.
<i>neutral RC</i> + <i>because</i>	He respected the waitress who was serving his table because over the past year she had received two promotions and was now practically in charge.
<i>neutral RC</i> + <i>and so</i>	He respected the waitress who was serving his table and so over the past year she had received very generous tips whenever he visited.
<i>causal RC</i> + <i>because</i>	He respected the waitress who could carry the heaviest trays because over the past year she had received two promotions and was now practically in charge.
<i>causal RC</i> + <i>and so</i>	He respected the waitress who could carry the heaviest trays and so over the past year she had received very generous tips whenever he visited.
Wrap-up	When his steak arrived, it was perfectly medium-rare.
<i>Verification statement</i>	<i>Mr. Evans' steak was overcooked.</i>

28.

Intro	Right when the gardening crew was leaving, Lady Paulson noticed that her beloved flower patch had been completely ruined.
<i>neutral RC</i> + <i>because</i>	She blamed the gardener who was wearing dark green overalls because surely he was the one last seen near the flower patch.
<i>neutral RC</i> + <i>and so</i>	She blamed the gardener who was wearing dark green overalls and so surely he was the one most fervently denying having been near the flowers.
<i>causal RC</i> + <i>because</i>	She blamed the gardener who always had a careless attitude because surely he was the one last seen near the flower patch.
<i>causal RC</i> + <i>and so</i>	She blamed the gardener who always had a careless attitude and so surely he was the one most fervently denying having been near the flowers.
Wrap-up	The crew manager promised they would repair the damage.
<i>Verification statement</i>	<i>Lady Paulson noticed her orchard had been ruined.</i>

29.

Intro	Julie was a contestant in a national talent show.
<i>neutral RC</i> + <i>because</i>	She envied the singer who was standing on stage left because from the start he had been the public's favourite.
<i>neutral RC</i> + <i>and so</i>	She envied the singer who was standing on stage left and so from the start he had been avoiding her as much as possible.
<i>causal RC</i> + <i>because</i>	She envied the singer who was also a gifted guitar player because from the start he had been the public's favourite.
<i>causal RC</i> + <i>and so</i>	She envied the singer who was also a gifted guitar player and so from the start he had been avoiding her as much as possible.
Wrap-up	She was extremely nervous about next week's final.

30.

Intro	Mr. Lee was helping the golf team off the bus after the accident.
<i>neutral RC</i> + <i>because</i>	He carried the girl who had been in the back of the bus because obviously she was very badly hurt.
<i>neutral RC</i> + <i>and so</i>	He carried the girl who had been in the back of the bus and so obviously she was very grateful for his help.
<i>causal RC</i> + <i>because</i>	He carried the girl who had a huge gash on her leg because obviously she was very badly hurt.
<i>causal RC</i> + <i>and so</i>	He carried the girl who had a huge gash on her leg and so obviously she was very grateful for his help.
Wrap-up	Fortunately, it would later turn out that the cut had not severed any major arteries.

31.

Intro	Billy was in his first year of primary school.
<i>neutral RC</i> + <i>because</i>	He feared the teacher who supervised the playground because every day she would come up to him and point out something he was doing wrong.
<i>neutral RC</i> + <i>and so</i>	He feared the teacher who supervised the playground and so every day she would come up to him and try to put him at ease.
<i>causal RC</i> + <i>because</i>	He feared the teacher who often yelled at students because every day she would come up to him and point out something he was doing wrong.
<i>causal RC</i> + <i>and so</i>	He feared the teacher who often yelled at students and so every day she would come up to him and try to put him at ease.
Wrap-up	He liked the PE teacher much better.

32.

Intro	Ava had to attend a symposium on life after high school.
<i>neutral RC</i> + <i>because</i>	She loudly mocked the speaker who presented before the break because right away he started to advocate expressing one's feelings using a xylophone.
<i>neutral RC</i> + <i>and so</i>	She loudly mocked the speaker who presented before the break and so right away he started to call her out on it.
<i>causal RC</i> + <i>because</i>	She loudly mocked the speaker who tripped as he walked on stage because right away he started to advocate expressing one's feelings by using a xylophone.
<i>causal RC</i> + <i>and so</i>	She loudly mocked the speaker who tripped as he walked on stage and so right away he started to call her out on it.
Wrap-up	The rest of the day was fairly disastrous also.

References

- Aijmer, Karin (1997). I think – An English modal particle. In T. Swan & O. Jansen Westvik (Eds.), *Modality in Germanic Languages: Historical and Comparative Perspectives* (pp. 1-48). Berlin & New York: Mouton de Gruyter.
- Andersson, Marta & Spenader, Jennifer (2014). Result and Purpose relations with and without ‘so’. *Lingua* 148, 1-27.
- Asher, Nicholas & Lascarides, Alex (2005). *Logics of Conversation*. Cambridge: Cambridge University Press.
- Asr, Fatimeh T. & Demberg, Vera (2012). Implicitness of discourse relations. In *Proceedings of the 24th International Conference on Computational Linguistics (COLING2012)* (pp. 2669-2684). Mumbai, India.
- Asr, Fatimeh T. & Demberg, Vera (2015). Uniform information density at the level of discourse relations: Negation markers and discourse connective omission. In *Proceedings of the 11th International Conference on Computational Semantics (IWCS2015)* (pp. 118-128). London, UK.
- Au, Terry K. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language* 25(1), 104-122.
- Baayen, R. Harald (2008) *Analyzing Linguistic Data: A Practical Introduction to Statistics Using R*. Cambridge: Cambridge University Press.
- Baayen, R. Harald, Davidson, Doug J. & Bates, Douglas M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59(4), 390-412.
- Bache, Carl & Jakobsen, Leif K. (1980). On the distinction between restrictive and non-restrictive relative clauses in Modern English. *Lingua* 52(3-4), 243-267.
- Bahdanau, Dzmitry, Cho, Kyunghyun, & Bengio, Yoshua (2014). *Neural machine translation by jointly learning to align and translate*. arXiv:1409.0473v7
- Barlow, Michael (2008). Parallel texts and corpus-based contrastive analysis. In M. De los Angeles Gómez-González, J. Lachlan McKenzie, & E.M. González Álvarez (Eds.), *Current Trends in Contrastive Linguistics. Functional and Cognitive Perspectives* (pp. 101-121). Amsterdam: John Benjamins.
- Barr, Dale J., Levy, Roger, Scheepers, Christoph, & Tily, Harry J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68(3), 255-278.
- Bates, Elisabeth A. (1976). *Language and Context: The Acquisition of Pragmatics*. New York: Academic Press.
- Bates, Douglas, Maechler, Martin, Bolker, Ben, Walker, Steve (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67(1), 1-48.

- Bazzanella, Carla (2011). Redundancy, repetition, and intensity in discourse. *Language Sciences* 33(2), 243-254.
- Becher, Victor (2011). When and why do translators add connectives? *Target* 23(1), 26-47.
- Bentivogli, Luisa, Bisazza, Arianna, Cettolo, Mauro, & Federico, Marcello (2016). Neural versus phrase-based Machine Translation quality: A case study. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 257-267). Austin, Texas.
- Biber, Douglas (2006). Stance in spoken and written university registers. *Journal of English for Academic Purposes* 5, 97-116.
- Biber, Douglas & Finegan, Edward (1989). Styles of stance in English: Lexical and grammatical marking of evidentiality and affect. *Text* 9(1), 93-124.
- Bisiada, Mario (2016). Lösen sie Schachtelsätze möglichst auf: The impact of editorial guidelines on sentence splitting in German business article translations. *Applied Linguistics* 37(3), 354-376.
- Blakemore, Diane (1987). *Semantic Constraints on Relevance*. Oxford: Blackwell.
- Bloom, Lois, Lahey, Margaret, Hood, Lois, Lifter, Karin, & Fiess, Kathleen (1980). Complex sentences: acquisition of syntactic connectives and the semantic relations they encode. *Journal of Child Language* 7, 235-261.
- Blum-Kulka, Shoshana (1986). Shifts of cohesion and coherence in translation. In J. House & S. Blum-Kulka (Eds.), *Interlingual and Intercultural Communication* (pp. 17-36). Tübingen: Gunter Narr Verlag.
- Brinton, Laurel J. (2008). *The Comment Clause in English*. Cambridge, Cambridge University Press.
- Canestrelli, Anneloes R. (2013). *Small Words, Big Effects? Subjective versus Objective Causal Connectives in Discourse Processing*. PhD thesis, Utrecht University. Utrecht: LOT. Available online at http://www.lotpublications.nl/Documents/325_fulltext.pdf
- Canestrelli, Anneloes R., Mak, Willem M. & Sanders, Ted J.M. (2013). Causal connectives in discourse processing: How differences in subjectivity are reflected in eye movements. *Language and Cognitive Processes* 28(9), 1394-1413.
- Carlson, Katy (2014). Predicting contrasts in sentences with and without focus marking. *Lingua* 150, 78-91.
- Carlson, Lynn & Marcu, Daniel (2001). *Discourse Tagging Reference Manual*. ISI technical report ISI-TR-545. doi: <ftp://128.9.176.20/isi-pubs/tr-545.pdf>
- Carlson, Lynn, Okurowski, Mary Ellen, & Marcu, Daniel (2002). *RST Discourse Treebank*. Web Download. Philadelphia: Linguistic Data Consortium.
- Carston, Robyn (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. New York, NY: John Wiley & Sons.

- Cartoni, Bruno, Zufferey, Sandrine, & Meyer, Thomas (2013). Using the Europarl corpus for cross-linguistic research. *Belgian Journal of Linguistics* 27(1), 23-42.
- Cartoni, Bruno, Zufferey, Sandrine, Meyer, Thomas, & Popescu-Belis, Andrei (2011). How comparable are parallel corpora? Measuring the distribution of general vocabulary and connectives. In *Proceedings of the 4th Workshop on Building and Using Comparable Corpora* (pp. 78-86). Portland, OR, USA.
- Clark, Eve V. (2003). *First Language Acquisition*. Cambridge: Cambridge University Press.
- Clark, Herb H. (1974). Semantics and comprehension. In Sebeok, T.A. (Ed.), *Current Trends in Linguistics, Vol. 12: Linguistics and Adjacent Arts and Sciences* (pp. 1291-1498). The Hague: Mouton.
- Clayman, Steven E. (2002). Sequence and solidarity. In S.R. Thye, & E. Lawler (Eds.), *Advances in Group Processes (Vol. 19)* (pp. 229-253). Bradford: Emerald Group Publishing Limited.
- Cohen, Jacob (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20(1), 37-46.
- Commandeur, Edwin (2010). *Implicit Causality and Implicit Consequentiality in Language Comprehension*. PhD Thesis, Tilburg University. Enschede: Gildeprint drukkerijen.
- Crible, Ludvine & Demberg, Vera (2018). *Disambiguating discourse relations with or without a connective: Does “and” really say nothing?* Poster presented at the Final TextLink Action Conference, Toulouse, France.
- Crinean, Marcelle & Garnham, Alan (2006). Implicit causality, implicit consequentiality and semantic roles. *Language and Cognitive Processes* 21(5), 636-648.
- Conrad, Susan & Biber, Douglas (2000). Adverbial marking of stance in speech and writing. In G. Thompson & S. Hunston (Eds.), *Evaluation in Text: Authorial Stance and the Construction of Discourse* (pp. 56-74). Oxford: Oxford University Press.
- Cornillie, Bert (2009). Evidentiality and epistemic modality: On the close relationship between two different categories. *Functions of Language* 16(1), 44-62.
- Dancygier, Barbara (1985). If, unless, and their Polish equivalents. *Papers and Studies in Contrastive Linguistics* 20, 65-72.
- Dancygier, Barbara & Sweetser, Eve E. (2005). *Mental Spaces in Grammar: Conditional Constructions*. Cambridge: Cambridge University Press.
- Dancygier, Barbara & Sweetser, Eve E. (2012). *Viewpoint in Language: A Multimodal Perspective*. Cambridge: Cambridge UP.
- Das, Debopam, & Taboada, Maite (2013). Explicit and implicit coherence relations: A corpus study. In *Proceedings of the 2013 Annual Conference of the Canadian Linguistic Association*. Victoria, Canada.

- Das, Debopam, Taboada, Maite, and McFetridge, Paul (2015). *RST Signalling Corpus LDC2015T10*. Web Download. Philadelphia: Linguistic Data Consortium.
- Degand, Liesbeth (1996). Causation in Dutch and French: interpersonal aspects. In R. Hasan, C. Cloran, & D.G. Butt (Eds.), *Functional Descriptions: Theory in Practice* (pp. 207-237). Amsterdam: John Benjamins.
- Degand, Liesbeth (2004). Contrastive analyses, translation and speaker involvement: The case of *puisque* and *aangezien*. In M. Achard & S. Kemmer (Eds.), *Language, Culture and Mind* (pp. 251-270). Stanford: CSLI Publications.
- Degand, Liesbeth & Pander Maat, Henk L.W. (2003). A contrastive study of Dutch and French causal connectives on the Speaker Involvement Scale. In A. Verhagen & J. van de Weijer (Eds.), *Usage Based Approaches to Dutch* (pp. 175-199). Utrecht: LOT.
- Degand, Liesbeth & Sanders, Ted J.M. (1999). Causal connectives in language use: Theoretical and methodological aspects of the classification of coherence relations and connectives. In *Levels of Representation in Discourse, Working Notes of the International Workshop on Text Representation* (pp. 3-12). Edinburgh, UK.
- Demberg, Vera, Asr, Fatimeh T., & Scholman, Merel C.J. (2017). *How compatible are our discourse annotations? Insights from mapping RST-DT and PDTB annotations*. arXiv:1704.08893v2
- Dupont, Maite, & Zufferey, Sandrine (2017). Methodological issues in the use of parallel directional corpora: A case study with English and French concessive connectives. *International Journal of Corpus Linguistics* 22(2), 270-297.
- Dyvik, Helge (1998). A translational basis for semantics. In S. Johansson & S. Oksefjell (Eds.), *Corpora and Cross-linguistic Research: Theory, Method, and Case Studies* (pp. 51-86). Amsterdam: Rodopi.
- Ehrlich, Kate (1980). Comprehension of pronouns. *Quarterly Journal of Experimental Psychology* 32(2), 247-255.
- Eisenberg, Ann R. (1980). A syntactic, semantic and pragmatic analysis of conjunction. *Stanford Papers and Reports on Child Language Development* 19, 70-78.
- Evers-Vermeul, Jacqueline (2005). *The Development of Dutch Connectives: Change and Acquisition as Windows on Form-Function Relations*. PhD thesis, Utrecht University. Utrecht: LOT. Available online at http://www.lotpublications.nl/Documents/110_fulltext.pdf
- Evers-Vermeul, Jacqueline, Hoek, Jet, & Scholman, Merel C.J. (2017). On temporality in discourse annotation: Theoretical and practical considerations. *Dialogue & Discourse* 8(2), 1-20.

- Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2009). The emergence of Dutch connectives; how cumulative cognitive complexity explains the order of acquisition. *Journal of Child Language* 36(4), 829-854.
- Evers-Vermeul, Jacqueline & Sanders, Ted J.M. (2011). Discovering domains – On the acquisition of causal connectives. *Journal of Pragmatics* 43(6), 1645-1662.
- Fabb, Nigel (1990). The difference between English restrictive and nonrestrictive relative clauses. *Linguistics* 26(1), 57-78.
- Fabricius-Hansen, Cathrine (1999). Information packaging and translation: Aspects of sentence splitting (German - English/Norwegian). In M. Doherty (Ed.), *Sprachspezifische Aspekte der Informationsverteilung* (pp. 175-214). Akademie-Verlag, Berlin.
- Fauconnier, Gilles (1985). *Mental Spaces*. Cambridge, MA: The MIT Press.
- Feinstein, Alvan R. & Cicchetti, Domenic V. (1990). High agreement but low kappa: I. The problem of two paradoxes. *Journal of Clinical Epidemiology* 43(6), 543-549.
- Ferstl, Evelyn C., Garnham, Alan, & Manouilidou, Christina (2011). Implicit causality bias in English: a corpus of 300 verbs. *Behavior Research Methods* 43(1), 124-135.
- Fetzer, Anita (2014). *I think, I mean and I believe* in political discourse: Collocates, functions and distribution. *Functions of Language* 21(1), 67-94.
- Fine, Alex B., Jaeger, T. Florian, Farmer, Thomas A., & Qian, Ting (2013). Rapid expectation adaptation during syntactic comprehension. *PloS ONE* 8(10).
- Ford, Cecelia E. (1993). *Grammar in Interaction: Adverbial Clauses in American English Conversations*. Cambridge: Cambridge University Press.
- Frank, Austin F. & Jaeger, T. Florian (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. In *Proceedings of the 28th Meeting of the Cognitive Science Society* (pp. 939-944). Washington DC, USA.
- Garvey, Catherine & Caramazza, Alfonso (1974). Implicit causality in verbs. *Linguistic Inquiry* 5(3), 459-464.
- Gamut, L.T.F. (1982). *Logica, Taal en Betekenis*. [Logic, Language and Meaning.] (Vol.2). Utrecht: Spectrum.
- Geeraerts, Dirk & Cuyckens, Hubert (2010). *The Oxford Handbook of Cognitive Linguistics*. Oxford: Oxford University Press.
- Givón, Talmy (1993). *English Grammar: A Function-Based Introduction*. Amsterdam: Benjamins.
- Granger, Sylviane & Tyson, Stephanie (1996). Connector usage in the English essay writing of native and non-native EFL speakers of English. *World Englishes* 15(1), 17-27.

- Grice, Paul (1975). Logic and conversation. In P. Cole and J.L. Morgan (Eds.), *Syntax and Semantics 3: Speech Acts* (pp. 41-58). New York: Academic Press.
- Grisot, Cristina & Blochowiak, Joanna (submitted). *Online processing of temporal discourse relations: Evidence from French using self-paced reading experiments*.
- Grosz, Barbara J. & Sidner, Candace L. (1986). Attention, intentions and the structure of discourse. *Computational Linguistics* 12, 175-204.
- Gwet, Kilem (2002). Kappa statistic is not satisfactory for assessing the extent of agreement between raters. *Statistical Methods for Inter-Rater Reliability Assessment* 1, 1-6.
- Haberlandt, Karl & Bingham, Geoffrey (1978). Verbs contribute to the coherence of brief narratives: Reading related and unrelated sentence triples. *Journal of Verbal Learning and Verbal Behavior* 17(4), 419-425.
- Halliday, Michael A.K. & Hasan, Ruqaiya (1976). *Cohesion in English*. London & New York: Routledge.
- Halverson, Sandra (2004). Connectives as a translation problem. In: H. Kittel et al. (Eds.), *An International Encyclopedia of Translation Studies* (pp. 562-572). Berlin/New York: Walter de Gruyter.
- Hansen-Schirra, Silvia, Neumann, Stella, & Steiner, Erich (2007). Cohesive explicitness and explicitation in an English-German translation corpus. *Languages in Contrast* 7(2), 241-265.
- Hartshorne, Joshua K. (2014). What is implicit causality? *Language, Cognition and Neuroscience* 29(7), 804-824.
- Hobbs, Jerry R. (1978). *Why is discourse coherent?* Technical Report 176, SRI International, Menlo Park, CA.
- Hobbs, Jerry R. (1979). Coherence and coreference. *Cognitive Science* 3, 67-90.
- Hobbs, Jerry R. (1990). *Literature and Cognition*. Stanford: CSLI.
- Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2015). The role of expectedness in the implicitation and explicitation of discourse relations. In *Proceedings of the Second Workshop on Discourse in Machine Translation (DiscoMT)* (pp. 41-46). Lisbon, Portugal.
- Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2016). *Discourse segmentation and ambiguity in discourse structure*. Poster presented at the Second TextLink Action Conference, Budapest, Hungary.
- Hoek, Jet & Scholman, Merel C.J. (2017). Evaluating discourse annotation: Some recent insights and new approaches. In *Proceedings of the 13th Joint ISO-ACL Workshop on Interoperable Semantic Annotation (isa-13)* (pp. 1-13). Montpellier, France.
- Hoek, Jet, Zufferey, Sandrine, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (in prep). *CCR-ParCor*.

- Hoey, Michael (1983). *On the Surface of Discourse*. London: George Allen & Unwin.
- Horn, Laurence R. (1984). Toward a new taxonomy for pragmatic inference. In D. Schiffrin (Ed.), *Meaning, Form, and Use in Context: Linguistic Applications* (pp. 11–42). Washington, D.C.: Georgetown University Press.
- Horn, Laurence R. (1993). Economy and redundancy in a dualistic model of natural language. *SKY 1993: 1993 Yearbook of the Linguistic Association of Finland*, 33-72.
- Hothorn, Torsten, Bretz, Frank, & Westfall, Peter (2008). Simultaneous Inference in General Parametric Models. *Biometrical Journal* 50(3), 346-363.
- Hovy, Eduard H. & Maier, Elisabeth (1995). *Parsimonious or profligate: How many and which discourse structure relations?* Unpublished manuscript.
- Hunter, Julie (2016). Reports in discourse. *Dialogue & Discourse* 7(4), 1-35.
- Isabelle, Pierre, Cherry, Colin, & Foster, George (2017). A challenge set approach to evaluating machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 2486-2496). Copenhagen, Denmark.
- Jordan, Michael P. (1984). *Rhetoric of Everyday English Texts*. London: George Allen & Unwin.
- Jordan, Michael P. (1988). The power of negation in English: Text, context and relevance. *Journal of Pragmatics* 29(6), 705-752.
- Junczys-Dowmunt, Marcin, Dwojak, Tomasz, & Hoang, Hieu (2016). Is Neural Machine Translation Ready for Deployment? A Case Study on 30 Translation Directions. In *Proceedings of the 13th International Workshop on Spoken Language Translation (IWSLT)*. Seattle, WA, USA.
- Kaltenböck, Gunther (2009). Initial *I think*: Main or comment clause? *Discourse and Interaction* 2(1), 49-70.
- Katz, Evelyn W. & Brent, Sandor D. (1968). Understanding connectives. *Journal of Verbal Learning and Verbal Behavior* 7, 501-509.
- Kärkkäinen, Elise (2003). *Epistemic Stance in English Conversation. A Description of Its Interactional Functions, with a Focus on 'I Think'*. Amsterdam: Benjamins.
- Keenan, Janice M., Baillet, Susan D., & Brown, Polly (1984). The effects of causal cohesion on comprehension and memory. *Journal of Verbal Learning and Verbal Behavior* 23(2), 115-126.
- Kehler, Andrew (1994). Temporal relations: reference or discourse coherence? In *Proceedings of the 32nd Annual Meeting on Association for Computational Linguistics* (pp. 319-321). Las Cruces, NM, USA.
- Kehler, Andrew (1995). *Interpreting Cohesive Forms in the Context of Discourse Inference*. PhD thesis, Harvard University.
- Kehler, Andrew (2002). *Coherence, Reference, and the Theory of Grammar*. Stanford: CSLI.

- Kehler, Andrew, Kertz, Laura, Rohde, Hannah, & Elman, Jeffrey L. (2008). Coherence and Coreference revisited. *Journal of Semantics* 25(1), 1-44.
- Kehler, Andrew & Rohde, Hannah (2015). Pronominal reference and pragmatic enrichment: A Bayesian account. *Proceedings of the 37th Annual Conference of the Cognitive Science Society (CogSci2015)* (pp. 1063-1068). Pasadena, CA, USA.
- Kintsch, Walter (1998). *Comprehension: A Paradigm for Cognition*. Cambridge: Cambridge University Press.
- Kleijn, Suzanne (2012). *On-line Processing of Objective and Subjective Causal Connectives*. MA thesis, Utrecht University. Available online at <https://dspace.library.uu.nl/bitstream/handle/1874/253962/S.%20Kleijn%20-%20On-line%20processing%20of%20subjective%20and%20objective%20causal%20connectives.pdf?sequence=2&isAllowed=y>
- Knott, Alistair (2001). Semantic and pragmatic relations and their intended effects. In T.J.M. Sanders, J. Schilperoord, & W.P.M.S. Spooren, (Eds.), *Text Representation: Linguistic and Psycholinguistic Aspects* (pp. 127-151). Amsterdam & Philadelphia: John Benjamins.
- Knott, Alistair & Dale, Robert (1994). Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes* 18(1), 35-62.
- Knott, Alistair & Sanders, Ted J.M. (1998). The classification of coherence relations and their linguistic markers: An exploration of two languages. *Journal of Pragmatics* 30(2), 135-175.
- Koehn, Philipp (2005). Europarl: A parallel corpus for statistical machine translation. In *Proceedings of the Tenth Machine Translation Summit (MT Summit X)*. Phuket, Thailand.
- Koehn, Philipp, Och, Franz J., & Marcu, Daniel (2003). Statistical phrase-based translation. In *Proceedings of HLT-NAACL2003* (pp. 48-54). Edmonton, Canada.
- Koehn, Philipp & Knowles, Rebecca (2017). Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation* (pp. 28-39). Vancouver, Canada.
- Köhne, Judith & Demberg, Vera (2013). The time-course processing of discourse connectives. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society (CogSci2013)* (pp. 2760-2765). Berlin, Germany.
- Koornneef, Arnout W. & Sanders, Ted J.M. (2013). Establishing coherence relations in discourse: The influence of implicit causality and connectives on pronoun resolution. *Language and Cognitive Processes* 28(8), 1169-1206.
- Koornneef, Arnout W. & van Berkum, Jos J.A. (2006). On the use of verb-based implicit causality in sentence comprehension: Evidence from eye tracking. *Journal of Memory and Language* 54(4), 445-465.

- Kortmann, Bernd (1991). *Free Adjuncts and Absolutes in English*. London: Routledge.
- Krippendorff, Klaus (2004). Reliability in content analysis: Some common misconceptions and recommendations. *Human Communication Research* 30(3), 411-433.
- Kruger, Haidee (2017). The effects of editorial intervention: Implications for studies of the features of translated language. In G. De Sutter, M. Lefer, & I. Delaere, (Eds.), *Empirical Translation Studies: New Methodological and Theoretical Traditions (Vol. 300 of TILSM series)* (pp. 113-156). The Hague: Mouton de Gruyter.
- Kuperberg, Gina R. & Jaeger, T. Florian (2016). What do we mean by prediction in language comprehension? *Language, Cognition, and Neuroscience* 31(1), 32-59.
- Kutas, Marta, DeLong, Katherine A., Smith, Nathaniel J. (2011). A look around at what lies ahead: Prediction and predictability in language processing. In M. Bar (Ed.), *Predictions in the Brain: Using Our Past to Generate a Future* (pp. 190-207). Oxford: Oxford University Press.
- Langacker, Ronald W. (1990). Subjectification. *Cognitive Linguistics* 1(1), 5-38.
- Langacker, Ronald W. (1991). Cognitive grammar. In F.G. Droste & J.E. Joseph (Eds.), *Linguistic Theory and Grammatical Description: Nine Current Approaches* (pp. 275-306). Amsterdam: John Benjamins.
- Langacker, Ronald W. (2006). Subjectification, grammaticization, and conceptual archetypes. In A. Athanasiadou, C. Canakis, & B. Cornillie (Eds.), *Subjectification* (pp. 17-40). Berlin: Mouton De Gruyter.
- Lascarides, Alex & Asher, Nicholas (1993). Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy* 16(5), 437-493.
- Levshina, Natalia & Degand, Liesbeth (2017). Just because: In search of objective criteria of subjectivity expressed by causal connectives. *Dialogue & Discourse* 8(1), 132-150.
- Levinson, Stephen C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Levy, Roger & Jaeger, T. Florian (2007). Speakers optimize information density through syntactic reduction. In B. Schlökopf, J. Platt, & T. Hoffman (Eds.), *Advances in Neural Information Processing Systems (NIPS) (Vol. 19)* (pp. 849-856). Cambridge, MA: MIT Press.
- Li, Fang (2014). *Subjectivity in Mandarin Chinese*. PhD thesis, Utrecht University. Utrecht: LOT. Available online at https://www.lotpublications.nl/Documents/365_fulltext.pdf

- Li, Fang, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2016). On the subjectivity of Mandarin reason connectives: Robust profiles or genre-sensitivity? In N.M. Stukker, W.P.M.S. Spooren, & G.J. Steen (Eds.), *Genre in Language, Discourse and Cognition* (pp. 15-49). Amsterdam: De Gruyter Mouton.
- Lison, Pierre & Tiedemann, Jörg (2016). *OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles*. Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016), 923-929. Portoroz, Slovenia.
- Longacre, Robert E. (1983). *The Grammar of Discourse*. New York: Plenum Press.
- Mak, Willem M., & Sanders, Ted J.M. (2013). The role of causality in discourse processing: Effects of expectation and coherence relations. *Language and Cognitive Processes*, 28(9), 1414-1437.
- Mann, William C. & Thompson, Sandra A. (1988). Rhetorical Structure Theory: Toward a functional theory of text organization. *Text* 8(3), 243-281.
- Mariño, José. Banchs, Rafael E., Crego, Josep M., & de Gispert, Adrià (2006). N-gram-based Machine Translation. *Computational Linguistics* 32(4), 527-549.
- Martin, James (1992). *English Text: Systems and Structure*. Amsterdam: John Benjamins.
- Matthiessen, Christian & Thompson, Sandra A. (1988). The structure of discourse and 'subordination.' In J. Haiman & S.A. Thompson (Eds.), *Clause Combining and Grammar in Discourse* (pp. 275-329). Amsterdam & Philadelphia: John Benjamins.
- McKoon, Gail, Greene, Steven B. & Ratcliff, Roger (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19(5), 1040-1052.
- McNamara, Danielle S., Kintsch, Eileen, Songer, Nancy B., & Kintsch, Walter (1996). Are good texts always better? Interactions of text coherence, background knowledge, and levels of understanding in learning from text. *Cognition and Instruction* 14(1), 1-43.
- Melamed, I. Dan (2001). *Empirical Methods for Exploiting Parallel Texts*. Cambridge, MA: MIT Press.
- Meyer, Thomas (2015). *Discourse-level Features for Statistical Machine Translation*. PhD thesis, École Polytechnique Fédérale de Lausanne.
- Meyer, Thomas, Hajlaoui, & Popescu-Belis (2015). Disambiguating discourse connectives for statistical machine translation. *IEEE Transactions on Audio, Speech, and Language Processing* 23(7), 1184-1197.

- Meyer, Thomas, & Popescu-Belis, Andrei (2012). Using sense-labeled discourse connectives for statistical machine translation. In *Proceedings of the Joint Workshop on Exploiting Synergies between Information Retrieval and Machine Translation (ESIRMT) and Hybrid Approaches to Machine Translation (HyTra)* (pp. 129-138). Avignon, France.
- Meyer, Thomas & Webber, Bonnie L. (2013). Implication of discourse connectives in (machine) translation. In *Proceedings of the Workshop on Discourse in Machine Translation (DiscoMT)* (pp. 19-26). Sofia, Bulgaria.
- Miculicich Werlen, Lesly, Pappas, Nikolaos, Ram, Dhananjay, & Popescu-Belis, Andrei (2018). Self-attentive residual decoder for neural machine translation. To appear at *NAACL-HLT 2018*. arXiv:1709.04849v4
- Miltsakaki, Eleni, Prasad, Rashmi, Joshi, Aravind K, & Webber, Bonnie L. (2004). The Penn Discourse Treebank. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)* (pp. 2237-2240). Lisbon, Portugal.
- Müller, Simone (2005). *Discourse Markers in Native and Non-native English Discourse*. Amsterdam/Philadelphia: John Benjamins.
- Murray, John D. (1997). Connectives and narrative text: The role of continuity. *Memory and Cognition* 25(2), 227-236.
- Myers, Jerome L., Shinjo, Mikiko, & Duffy, Susan A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language* 26(4), 453-465.
- Noël, Dirk (2003). Translations as evidence for semantics: an illustration. *Linguistics* 41(4), 757-785.
- Noordman Leo G.M. & de Blijzer, Femke (2000). On the processing of causal relations. In E. Couper-Kuhlen & B. Kortmann (Eds.), *Cause, Condition, Concession and Contrast: Cognitive and Discourse Perspectives* (pp. 35-56). Berlin/New York: Mouton de Gruyter.
- Pander Maat, Henk L.W. (1998). Classifying negative coherence relations on the basis of linguistic evidence. *Journal of Pragmatics* 30(2), 177-204.
- Pander Maat, Henk L.W. (2002). *Tekstanalyse [Text Analysis]*. Bussum: Coutinho.
- Pander Maat, Henk L.W. & Degand, Liesbeth (2001). Scaling causal relations and connectives in terms of speaker involvement. *Cognitive Linguistics* 12(3), 211-246.
- Pander Maat, Henk L.W. & Sanders, Ted J.M. (2000). Domains of use or subjectivity? The distribution of three Dutch causal connectives explained. *Topics in English Linguistics* 33, 57-82.
- Pander Maat Henk L.W. & Sanders, Ted J.M. (2001). Subjectivity in causal connectives: An empirical study of language in use. *Cognitive Linguistics* 12(2), 247-274.

- Pareti, Sylvia (2012). A database of attribution relations. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)* (pp. 3213-3217). Istanbul, Turkey.
- Patterson, Gary & Kehler, Andrew (2013). Predicting the presence of discourse connectives. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 914-923). Seattle, WA, USA.
- PDTB Research Group (2007). *The Penn Discourse Treebank 2.0 Annotation Manual*. IRCS technical report.
- Piaget, Jean (1969). *Judgement and Reasoning in the Child*. London: Routledge & Kegan Paul. [Translation of: Jean Piaget, 1924. *Le Jugement et le Raisonnement chez l'Enfant*. Neuchatel: Delachaux et Niestlé.]
- Pickering, Martin J. & Majid, Asifa (2007). What are implicit causality and consequentality? *Language and Cognitive Processes* 22(5), 780-788.
- Pit, Mirna (2003). *How to Express Yourself with a Causal Connective: Subjectivity and Causal Connectives in Dutch, German, and French*. Amsterdam: Rodopi.
- Prasad, Rashmi, Dinesh, Nikhil, Lee, Alan, Miltsakaki, Eleni, Robaldo, Livio, Joshi, Avarind, & Webber, Bonnie (2008). The Penn Discourse Treebank 2.0. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008)* (pp. 2961-2968). Marrakech, Morocco.
- Prasad, Rashmi, Webber, Bonnie L., Lee, Alan, & Joshi, Avarind (in prep). *Discourse relations in the PDTB 3.0*.
- Pomerantz, Anita & Heritage, John C. (2013). Preference. In J. Sidnell & T. Stivers (Eds.), *The Handbook of Conversation Analysis* (pp. 210-228). Chichester: John Wiley & Sons.
- Polanyi, Livia (1988). A formal model of the structure of discourse. *Journal of Pragmatics* 12(5-6), 601-638.
- R Core Team (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Redeker, Gisela (1990). Ideational and pragmatic markers of discourse structure. *Journal of Pragmatics* 14(3), 367-381.
- Reese, Brian, Hunter, Julie, Asher, Nicholas, Denis, Pascal, & Baldridge, Jason (2007). *Reference manual for the analysis of rhetorical structure*. Unpublished manuscript. University of Texas at Austin, Austin, TX. http://timeml.org/jamesp/annotation_manual.pdf
- Rehbein, Ines, Scholman, Merel C.J. & Demberg, Vera (2016). Annotating discourse relations in spoken language: A comparison of the PDTB and CCR frameworks. *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)* (pp. 1039--1046). Portoroz, Slovenia.

- Reid, James M. (2016). *Free adjuncts and temporal coherence*. Unpublished manuscript.
- Renkema, Jan (2004). *Discourse Studies: An Introductory Textbook*. Amsterdam: Benjamins.
- Renkema, Jan (2009). *The Texture of Discourse: Towards an outline of Connectivity Theory*. Amsterdam & Philadelphia: John Benjamins.
- Robaldo, Livio, & Miltsakaki, Eleni (2014). Corpus-driven semantics of concession: Where do expectations come from? *Dialogue & Discourse* 5(1), 1-36.
- Rohde, Hannah, Dickinson, Anna, Schneider, Nathan, Clark, Christopher N.L., Louis, Annie, & Webber, Bonnie L. (2016). Filling in the blanks in understanding discourse adverbials: Consistency, conflict, and context-dependence in a crowdsourced elicitation task. In *Proceedings of 10th Linguistic Annotation Workshop (LAW X)*. Berlin, Germany.
- Rohde, Hannah & Kehler, Andrew (2014). Grammatical and information-structural influences on pronoun production. *Cognition and Neuroscience* 29(8), 912-927.
- Rohde, Hannah, Kehler, Andrew, & Elman, Jeffrey L. (2006). Event structure and discourse coherence biases in pronoun interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci2006)* (pp. 697-702). Vancouver, Canada.
- Rohde, Hannah, Levy, Roger, & Kehler, Andrew (2011). Anticipating explanations in relative clause processing. *Cognition* 118(3), 339-358.
- Sanders, José (1994). *Perspectives in Narrative Discourse*. PhD thesis, Tilburg University. Available online at <https://pure.uvt.nl/portal/files/896522/305833.pdf>
- Sanders, José & Redeker, Gisela (1996). Speech and thought in narrative discourse. In G. Fauconnier & E.E. Sweetser (Eds.), *Spaces, Worlds and Grammar* (pp. 290-317). Chicago: University of Chicago Press.
- Sanders, José, Sanders, Ted J.M., & Sweetser Eve E. (2012). Responsible subjects and discourse causality: How mental spaces and perspective help identifying subjectivity in Dutch backward causal connectives. *Journal of Pragmatics* 44(2), 191-213.
- Sanders, José & Spooren, Wilbert P.M.S. (1997). Perspective, subjectivity, and modality from a cognitive linguistic point of view. In W. Liebert, G. Redeker, & L.R. Waugh (Eds.), *Discourse and Perspective in Cognitive Linguistics* (pp. 85-112). Amsterdam: John Benjamins.
- Sanders, Ted J.M. (1997). Semantic and pragmatic sources of coherence: On the categorization of coherence relations in context. *Discourse Processes* 24(1), 119-147.

- Sanders, Ted J.M. (2005). Coherence, causality and cognitive complexity in discourse. In *Proceedings/Actes SEM-05, First International Symposium on the Exploration and Modelling of Meaning* (pp. 105-114). Biarritz, France.
- Sanders, Ted J.M., Demberg, Vera, Hoek, Jet, Scholman, Merel C., Asr, Fatimeh T., Zufferey, Sandrine, & Evers-Vermeul, Jacqueline (in press). Unifying dimensions in discourse relations: How various annotation frameworks are related. *Corpus Linguistics and Linguistic Theory*.
- Sanders, Ted J.M. & Evers-Vermeul, Jacqueline (in press). Subjectivity and causality in discourse and cognition; Evidence from corpus analyses, acquisition and processing. In Ó. Loureda, I.M. Recio Fernández, L. Nadal, and A. Cruz (Eds.), *Methodological Approaches to Discourse Markers*. Amsterdam/Philadelphia: John Benjamins.
- Sanders, Ted J.M., & Noordman, Leo G.M. (2000). The role of coherence relations and their linguistic markers in text processing. *Discourse Processes* 29(1), 37-60.
- Sanders, Ted J.M., Sanders, José, Sweetser, Eve E. (2009). Causality, cognition and communication: A mental space analysis of subjectivity in causal connectives. In T.J.M. Sanders & E.E. Sweetser (Eds.), *Causal Categories in Discourse and Cognition* (pp. 19-60). Berlin: Mouton De Gruyter.
- Sanders, Ted J.M. & Spooren, Wilbert P.M.S. (1999). Communicative intentions and coherence relations. In W. Bublitz, U. Lenk, & E. Ventola (Eds.), *Coherence in Spoken and Written Discourse: How to Create It and How to Describe It* (pp. 235-250). Amsterdam/Philadelphia: John Benjamins.
- Sanders, Ted J.M. & Spooren, Wilbert P.M.S. (2007). Discourse and text structure. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford Handbook of Cognitive Linguistics* (pp. 916-941). Oxford: Oxford University Press.
- Sanders, Ted J.M. & Spooren, Wilbert P.M.S. (2009). Causal categories in discourse: Converging evidence from language use. In T.J.M. Sanders & E.E. Sweetser (Eds.), *Causal Categories in Discourse and Cognition* (pp. 205-246). Berlin: Mouton De Gruyter.
- Sanders, Ted J.M. & Spooren, Wilbert P.M.S. (2013). Exceptions to rules: a qualitative analysis of backward causal connectives in Dutch naturalistic discourse. *Text & Talk* 33(3), 377-398.
- Sanders, Ted J.M. & Spooren, Wilbert P.M.S. (2015). Causality and subjectivity in discourse: The meaning and use of causal connectives in spontaneous conversation, chat interactions and written text. *Linguistics* 53(1), 53-92.
- Sanders, Ted J.M., Spooren, Wilbert P.M.S., & Noordman, Leo G.M. (1992). Toward a taxonomy of coherence relations. *Discourse Processes* 15(1), 1-35.

- Sanders, Ted J.M., Spooren, Wilbert P.M.S. & Noordman, Leo G.M. (1993). Coherence relations in a cognitive theory of discourse representation. *Cognitive Linguistics* 4(2), 93-133.
- Sanders, Ted J.M. & van Wijk, Carel H. (1996). PISA—A procedure for analyzing the structure of explanatory texts. *Text* 16(1), 91-132.
- Santana, Andrea, Spooren, Wilbert P.M.S., Nieuwenhuijsen, Dorien, & Sanders, Ted J.M. (submitted). *Subjectivity in Spanish discourse: Explicit and implicit causal relations in different contexts*.
- Schilperoord, Joost & Verhagen, Arie (1998). Conceptual dependency and the clausal structure of discourse. In J. Koenig (Ed.), *Discourse and Cognition: Bridging the Gap* (pp. 141-163). Stanford, CA: CSLI Publications.
- Scholman, Merel C.J. (in prep). *Justifying distinctions between coherence relations*.
- Scholman, Merel C.J. & Demberg, Vera (2017). Examples and specifications that prove a point: Identifying elaborative and argumentative discourse relations. *Dialogue & Discourse* 8(2), 56-83.
- Scholman, Merel C.J., Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2016). A step-wise approach to discourse annotation: Towards a reliable categorization of coherence relations. *Dialogue & Discourse* 7(2), 1-28.
- Segal, Erwin M. & Duchan, Judith F. (1997). Interclausal connectives as indicators of structuring in narrative. In J. Costermans & M. Fayol (Eds.), *Processing Interclausal Relationships: Studies in the Production and Comprehension of Text* (pp. 95-119). Mahwah, NJ: Lawrence Erlbaum Associates.
- Segal, Erwin M., Duchan, Judith F., & Scott, Paula J. (1991). The role of interclausal connectives in narrative structuring: Evidence from adults' interpretations of simple stories. *Discourse Processes* 14(1), 27-54.
- Solstad, Torgrim & Bott, Oliver (2013). Towards a formal theory of explanatory biases in discourse. In *Proceedings of the 19th Amsterdam colloquium* (pp. 203-210). Amsterdam, The Netherlands.
- Sperber, Dan & Wilson, Deirdre (1985). Loose talk. In *Proceedings of the Aristotelian Society LXXXVI* (pp. 540-549). London, UK
- Spooren, Wilbert P.M.S. (1989). *Some Aspects of the Form and Interpretation of Global Contrastive Coherence Relations*. PhD Thesis, K.U. Nijmegen.
- Spooren, Wilbert P.M.S. (1997). The processing of underspecified coherence relations. *Discourse Processes* 24(1), 149-168.
- Spooren, Wilbert P.M.S. & Degand, Liesbeth (2010). Coding coherence relations: Reliability and validity. *Corpus Linguistics and Linguistic Theory* 6(2), 241-266.
- Spooren, Wilbert P.M.S. & Sanders, Ted J.M. (2008). The acquisition of coherence relations: On cognitive complexity in discourse. *Journal of Pragmatics* 40(12), 2003-2026.

- Staub, Adrian & Clifton, Charles (2006). Syntactic prediction in language comprehension: Evidence from either...or. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32(2), 425-436.
- Stevenson, Rosemary J., Crawley, Rosalind A., & Kleinman, David (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes* 9(4), 519-548.
- Stewart, Andrew J., Pickering, Martin J., & Sanford, Anthony J. (2000). The time course of the influence of implicit causality information: Focusing versus integration accounts. *Journal of Memory and Language* 42(3), 423-443.
- Stukker, Ninke M. (2005). *Causality Marking across Levels of Language Structure*. PhD thesis, Utrecht University. Utrecht, LOT. Available online at https://www.lotpublications.nl/Documents/118_fulltext.pdf
- Stukker, Ninke M. & Sanders, Ted J.M. (2012). Subjectivity and prototype structure in causal connectives: A cross-linguistic perspective. *Journal of Pragmatics* 44(2), 169-190.
- Stukker, Ninke M., Sanders, Ted J.M., & Verhagen, Arie (2008). Causality in verbs and in discourse connectives. Converging evidence of cross-level parallels in Dutch linguistic categorization. *Journal of Pragmatics* 40(7), 1296-1322.
- Sweetser, Eve E. (1990). *From Etymology to Pragmatics: The Mind-Body Metaphor in Semantic Structure and Semantic Change*. Cambridge: Cambridge University Press.
- Taboada, Maite (2006). Discourse markers as signals (or not) of rhetorical relations. *Journal of Pragmatics* 38(4), 567-592.
- Taboada, Maite & Das, Debopam (2013). Annotation upon annotation: Adding signalling information to a corpus of discourse relations. *Dialogue and Discourse* 4(2), 249-281.
- Taboada, Maite & Mann, William C. (2006). Rhetorical Structure Theory: Looking back and moving ahead. *Discourse Studies* 8(3), 423-459.
- Teubert, Wolfgang (1999). Corpus linguistics: A partisan view. *TELRI Newsletter* 8, 4-19.
- Thompson, Sandra A. (2002). "Object complements" and conversation: Towards a realistic account. *Studies in Language* 26(1), 125-163.
- Thompson, Sandra A. & Mulac, Anthony (1991). The discourse conditions for the use of the complementizer that in conversational English. *Journal of Pragmatics* 15(3), 237-251.
- Trabasso, Tom & van den Broek, Paul (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language* 24(5), 612-630.
- Traugott, Elizabeth C. (1995). Subjectification in grammaticalization. In D. Stein & S. Wright (Eds.), *Subjectivity and Subjectivisation* (pp. 31-54). Amsterdam: Benjamins.

- Traxler, Matthew, J., Bybee, Michael D., & Pickering, Martin J. (1997a). Influence of connectives on language comprehension: Eye-tracking evidence for incremental interpretation. *The Quarterly Journal of Experimental Psychology: Section A* 50(3), 481-497.
- Traxler, Matthew, J. & Gernsbacher, Morton A. (2011). *Handbook of Psycholinguistics*. Amsterdam: Elsevier.
- Traxler, Matthew J., Sanford, Anthony J., Aked, Joy P., & Moxey, Linda M., (1997b). Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 23(1), 88-101.
- Trueswell, John C., Tanenhaus, Michael K., & Garnsey, Susan M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language* 33(3), 285-318.
- Umbach, Carla (2005). Contrast and information structure: A focus-based analysis of *but*. *Linguistics* 43(1), 207-232.
- van den Broek, Paul (1990). Causal inferences in the comprehension of narrative texts. In A.C. Graesser & G.H. Bower (Eds.), *The Psychology of Learning and Motivation: Inferences and Text Comprehension* (pp. 175-194). San Diego, CA: Academic Press.
- van Dijk, Teun A. (1977). *Text and Context*. New York: Longman.
- van Silfhout, Gerdineke, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M. (2015). Connectives as processing signals: How students benefit in processing narrative and expository texts. *Discourse Processes* 52(1), 47-76.
- van Veen, Rosie (2011). *The Acquisition of Causal Connectives: The Role of Parental Input and Cognitive Complexity*. PhD thesis, Utrecht University. Utrecht: LOT. Available online at http://www.lotpublications.nl/Documents/286_fulltext.pdf
- Verhagen, Arie (2000). Concession implies causality, though in some other space. In E. Couper-Kuhlen & B. Kortmann (Eds.), *Cause, Condition, Concession and Contrast: Cognitive and Discourse Perspectives* (pp. 361-380). Berlin/New York: Mouton de Gruyter.
- Verhagen, Arie (2001). Subordination and discourse segmentation revisited, or: Why matrix clauses may be more dependent than complements. In T.J.M. Sanders, J. Schilperoord, & W.P.M.S. Spooren (Eds.), *Text Representation: Linguistic and Psycholinguistic Aspects* (pp. 337-357). Amsterdam: John Benjamins.
- Verhagen, Arie (2005). *Constructions of Intersubjectivity: Discourse, Syntax, and Cognition*. Oxford: Oxford University Press.
- Versley, Yannick & Gastel, Anna (2013). Linguistic tests for discourse relations in the TüBa-D/Z corpus of written German. *Dialogue and Discourse* 4(2), 142-173.

- Wason, Peter C. & Johnson-Laird, Philip (1972). *Psychology of Reasoning: Structure and Content*. Cambridge, MA: Harvard University Press.
- Webber, Bonnie L. (2013). What excludes an alternative in coherence relations? In *Proceedings of the 10th International Workshop on Computational Semantics (IWCS2013)* (pp. 276-287). Potsdam, Germany.
- Wei, Yipu (2018). *Causal Connectives and Perspective Markers in Chinese: The Encoding and Processing of Subjectivity in Discourse*. PhD thesis, Utrecht University. Utrecht: LOT. Available online at https://www.lotpublications.nl/Documents/482_fulltext.pdf
- Wilson, Deirdre & Sperber, Dan (1993). Linguistic form and relevance. *Lingua* 90(1-2), 1-25.
- Wilson, Deirdre & Sperber, Dan (2005). Relevance Theory. In L. Horn & G. Ward (Eds.), *The Handbook of Pragmatics* (pp. 607-632). New York: Wiley.
- Wolf, Florian & Gibson, Edward (2005). Representing discourse coherence: A corpus-based study. *Computational Linguistics* 31(2), 249-287.
- Wolfe, Michael B.W., Magliano, Joseph P., & Larsen, Benjamin (2005). Causal and semantic relatedness in discourse understanding and representation. *Discourse Processes* 39(2-3), 165-187.
- Xiao, Hongling Li, Fang, Sanders, Ted J.M., & Spooren, Wilbert P.M.S. (submitted). *The subjectivity of Mandarin reason connectives: A corpus study of written, spoken and microblog discourse*.
- Zhao, Xinshu, Liu, Jun S., & Deng, Ke (2013). Assumptions behind intercoder reliability indices. In C.T. Salmon (Ed.), *Communication Yearbook 36* (pp. 419-480). New York: Routledge.
- Ziemski, Michal, Junczys-Dowmunt, Marcin, & Pouliquen, Bruno (2016). The United Nations Parallel Corpus v1.0. LREC. 2016. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)* (pp. 3530-3534). Portoroz, Slovenia.
- Zufferey, Sandrine (2010). *Lexical Pragmatics and Theory of Mind*. Amsterdam: John Benjamins.
- Zufferey, Sandrine, & Cartoni, Bruno (2014). A multifactorial analysis of explicitation of translation. *Target* 26(3), 23-42.
- Zufferey, Sandrine & Degand, Liesbeth (2017). Annotating the meaning of discourse connectives in multilingual corpora. *Corpus Linguistics and Linguistic Theory* 13(2), 399-422.
- Zufferey, Sandrine, & Gygax, Pascal M. (2015). The role of perspective shifts for processing and translating discourse relations. *Discourse Processes* 53(7), 532-555.

- Zufferey, Sandrine, & Gygax, Pascal M. (2017). Processing connectives with a complex form-function mapping in L2: The case of French “en effet.” *Frontiers in Psychology: Language* 8, 1-11.
- Zwaan, Rolf A., Langston, Mark C., & Graesser, Arthur C. (1995). The construction of situation models in narrative comprehension. *Psychological Science* 6(5), 292-297.
- Zwaan, Rolf A. & Radvansky, Gabriel (1998). Situation models in language comprehension and memory. *Psychological Bulletin* 123, 162-185.

Nederlandse samenvatting

Introductie

Een tekst is geen verzameling losse zinnen, maar een coherent geheel waarin alle zinnen met elkaar in verbinding staan. Om een tekst goed te begrijpen, moet een lezer of luisteraar dus vaststellen hoe elke zin zich verhoudt tot de omringende tekst. Soms wordt dit proces versimpeld door de aanwezigheid van connectieven (*omdat, maar, immers, etc.*), verbindingswoorden die expliciet aangeven hoe een zin aan de voorgaande zin(nen) ervoor of erna gekoppeld moet worden om tot een coherente mentale representatie van de tekst te komen. In veel gevallen moet de lezer of luisteraar echter zelf afleiden welke coherentierelatie gelegd moet worden tussen twee zinnen in een discourse, bijvoorbeeld een oorzaak-gevolg-relatie, een contrast of een regel met een uitzondering. In voorbeeld (1) is het zaak om vast te stellen dat de handelingen die beschreven worden in de individuele zinnen moeten worden uitgevoerd in de volgorde waarin ze in de tekst genoemd worden. Dit wordt slechts één keer expliciet aangegeven, door het dikgedrukte *daarna* in de vierde zin. Het fragment bevat nog een connectief: *en*. Ook de stap omschreven na *en* moet worden uitgevoerd na het afronden van de voorgaande stap, maar de lezer kan dit niet afleiden uit het connectief *en*, dat enkel aangeeft dat er nog een stap aan de instructies wordt toegevoegd.

- (1) Schenk de koude koffie in een grote kan. Voeg het roomijs en de suiker toe. Mix alles goed met een staafmixer of blender. Voeg **daarna** de melk, de koffiemelk en een snufje kaneel toe **en** mix alles nogmaals goed door.¹

Hoewel de meeste coherentierelaties in (1) dus zonder (volledige) hulp van connectieven geïnfereerd moeten worden, zal dit fragment voor de meeste mensen volstrekt duidelijk zijn. Bij het fragment in (2) ligt dit anders.

- (2a) **Hoewel** het lastig is om alcohol te bevriezen, zijn gintonic-ijsjes niet onmogelijk om te maken.²
- (2b) Het is lastig om alcohol te bevriezen. Gintonic-ijsjes zijn niet onmogelijk om te maken.

Zonder *hoewel* wordt het fragment in (2a) moeilijk te begrijpen, zoals (2b) illustreert. De intuïtie dat sommige relaties makkelijker uit te drukken zijn zonder connectief dan

¹ <https://www.24kitchen.nl/recepten/ijskoffie>

² <https://www.metronieuws.nl/lifestyle/food/2018/04/drie-makkelijke-ijsjes-om-zelf-te-maken>

andere is al bevestigd op basis van corpusstudies, maar een duidelijke verklaring hiervoor is er nog niet. De eerste centrale vraag in dit proefschrift luidt daarom: Wanneer moet een coherentierelatie expliciet gemarkeerd worden door middel van een connectief, en wanneer kan een relatie impliciet kan blijven?

Lezers of luisteraars moeten niet alleen het juiste type coherentierelatie infereren, maar moeten ook bepalen welke twee segmenten in die relatie figureren. Zo drukt de derde zin in (3), *daardoor bevatten ze minder calorieën*, een resultaat uit van de tweede zin. Samen vormen zin twee en drie een motivatie voor de claim in de eerste zin. De laatste zin staat in vergelijking of zelfs contrast met de eerste zin: goede ijsjes versus slechte ijsjes.

- (3) Waterijsjes zijn de beste keus. Hier zit namelijk geen vet in en vaak ook minder suiker dan in chocolade-, room- of yoghurtijs. Daardoor bevatten ze minder calorieën. Ijsjes die minder goed scoren zijn extra grote waterijsjes en ijsjes met room of yoghurt.³

Waar in (1) elke zin een apart segment is en elk segment aan de zin ervoor verbonden kan worden, is de segmentatie van (3) dus veel complexer. De tweede hoofdvraag van dit proefschrift luidt daarom: Tussen welke delen van een discourse leggen mensen coherentierelaties, oftewel: welke delen van een tekst kunnen als discourse-segment fungeren?

Discourse-segmentatie

In onderzoek naar coherentierelaties wordt door middel van *discourse-segmentatie* aangegeven tussen welke delen van een tekst coherentierelaties te infereren zijn. Het uitgangspunt van Hoofdstuk 2 is dat in een goede discourse-segmentatie alle segmenten corresponderen met de informatie-eenheden ('idea units') in de mentale representatie van een discourse. Aan de hand van corpusvoorbeelden wordt beargumenteerd dat sommige bestaande, algemeen gebruikte segmentatieregels (afkomstig van Mann & Thompson 1988) niet altijd leiden tot een goede segmentatie.

De eerste regel die besproken wordt, is het strikt gescheiden houden van segmentatie en annotatie, het proces waarin de coherentierelaties in een discourse voorzien worden van een relatielabel dat omschrijft welk verband er tussen twee segmenten gelegd wordt (bijvoorbeeld oorzaak-gevolg, contrast). In de praktijk houdt deze regel in dat de discourse-segmenten geïdentificeerd worden op basis van syntactische kenmerken, zonder dat de interpretatie van een tekst een rol speelt. Een veelgebruikt criterium voor segmenten is dat ze minimaal een deelzin ('clause') zijn. Ook moeten segmenten doorgaans syntactisch en conceptueel onafhankelijk zijn, wat

³ <https://www.consumentenbond.nl/kindervoeding/kinderijsjes>

ertoe leidt dat bijvoorbeeld beperkende bijvoeglijke bijzinnen, beperkende betrekkelijke bijzinnen, en complementszinnen niet als op zichzelf staande discourse-segmenten aangemerkt worden; deze worden gezien als onderdeel van hun matrix-zin.

In Hoofdstuk 2 wordt beargumenteerd dat de interpretatie van een fragment een cruciale voorwaarde kan zijn voor een goede segmentatie van een coherentierelatie. Zo is het bijvoorbeeld onmogelijk om het verschil tussen (4) en (5) aan te duiden als de betekenis van de fragmenten buiten beschouwing gelaten moet worden. Beide voorbeelden bevatten een lijdend voorwerp met een beperkende betrekkelijke bijzin, gevolgd door het connectief *omdat* en nog een deelzin. Hoewel de oppervlaktestructuur voor beide voorbeelden dezelfde is, hebben ze een andere discourse-structuur. In (4) geeft de deelzin na *omdat* een reden voor alles wat ervoor komt, zoals de vierkante haken aangeven. In (5) geeft de deelzin na *omdat* alleen een reden voor *die sporten*, en de *omdat*-zin moet dus worden beschouwd als onderdeel van de beperkende betrekkelijke bijzin.

- (4) [MicroFiber is zeer geschikt voor mensen die sporten]_{S1} omdat [het weinig of geheel geen vocht aantrekt.]_{S2}⁴
- (5) [Ik ken veel mensen [die sporten]_{S1a} omdat [het “moet.”]_{S1b}]_{S1}⁵

Als de interpretatie van de fragmenten niet in acht genomen zou worden, zou (5) dezelfde segmentatie krijgen als (4). Op die manier worden de informatie-eenheden die *omdat* met elkaar verbindt dus onnauwkeurig weergegeven.

Een andere regel (ook voorgesteld door Mann & Thompson 1988) stelt dat *alle* onderdelen van een tekst opgenomen moeten worden in de segmentatie van de tekst. In Hoofdstuk 2 wordt beargumenteerd dat alle onderdelen die deel uitmaken van de propositionele inhoud van een tekst terug moeten komen in de segmentatie van die tekst, maar dat dit voor andere onderdelen, bijvoorbeeld voor *stance markers* (bepalingen van modaliteit), optioneel zou moeten zijn. De fragmenten in (6) en (7) bevatten bijvoorbeeld beide de *stance marker* *waarschijnlijk*. In (6) is *waarschijnlijk* echt onderdeel van het eerste segment dat door *omdat* aan het tweede segment verbonden wordt; dat Jan Steen een bierbrouwerij had aan de overkant van de gracht leidt tot de claim dat hij de man waarschijnlijk goed kende. In (7) staat *waarschijnlijk* ook voor *omdat*, maar maakt het geen deel uit van het eerste segment dat in de mentale representatie van het fragment aan het tweede segment wordt verbonden. Dat een passagier het gevoel had dat je wat extra's deed, leidt niet tot de claim dat je dit compliment hebt gekregen; een betere parafrase is “het is waarschijnlijk dat je dit compliment hebt gekregen omdat een passagier het gevoel had dat je wat extra's

⁴ <https://www.come-markt.nl/microfiber-.html>

⁵ <https://zwangerpuurgezond.nl/logboekthecla/theclas-kraamweken-hoe-hersteld-buik-bevalling/>

deed.” In (7) modificeert *waarschijnlijk* dus de hele relatie en niet slechts het eerste segment.

- (6) Jan Steen kende hem **waarschijnlijk** goed omdat hij zelf een bierbrouwerij had aan de overkant van de gracht.⁶
- (7) Dit compliment heb je **waarschijnlijk** gekregen omdat een passagier het gevoel had dat je wat extra’s deed.⁷

Bij nauwkeurige segmentatie van de fragmenten in (6) en (7) zou *waarschijnlijk* in (6) in S1 opgenomen moeten worden, maar *waarschijnlijk* in (7) buiten S1 gelaten moeten worden. Om dit verschil te bepalen is, net als bij (4) en (5), de betekenis van de fragmenten nodig. In Hoofdstuk 2 wordt niet alleen beargumenteerd dat adverbiale *stance markers* zoals *waarschijnlijk* slechts opgenomen zouden moeten worden in de segmentatie van een tekst als dat leidt tot nauwkeurige segmenten, maar wordt dit ook geopperd voor Engelse matrix-zinnen die een *stance marker* functie vervullen, zoals *I think (that)* ‘Ik denk dat’ en *I believe (that)* ‘Ik geloof dat.’⁸

De vraag tussen welke delen van een discourse mensen coherentierelaties infereren, wordt verder opgepakt in Hoofdstuk 6. In Hoofdstuk 2 wordt besproken dat beperkende betrekkelijke bijzinnen doorgaans een uitzondering vormen op de deelzin-als-segment-vuistregel, omdat beperkende betrekkelijke bijzinnen niet in verbinding staan met een andere deelzin, maar met een zelfstandig naamwoord, en omdat ze zowel syntactisch als conceptueel geïntegreerd zijn in hun matrix-zin. In het parallelle corpus dat gebruikt wordt in Hoofdstuk 4 en 5 zijn echter verschillende voorbeelden te vinden van constructies met een beperkende betrekkelijke bijzin die met een connectief vertaald zijn en andersom: coherentierelaties met een connectief die vertaald zijn door middel van een beperkende betrekkelijke bijzinsconstructie, zoals in voorbeeld (7). Dit wekt de indruk dat het wellicht toch mogelijk is een coherentierelatie te infereren tussen beperkende betrekkelijke bijzinnen en hun matrix-zin.

⁶ <https://www.volkskrant.nl/cultuur-media/rijksmuseum-koopt-topstuk-jan-steen-bd451e7b/>

⁷ <https://www.uber.com/nl/drive/partner-app/compliments/>

⁸ Toekomstig onderzoek zal moeten uitwijzen of in andere talen matrix-zinnen ook een *stance marker* functie kunnen vervullen.

- (23) EN Recently we have seen headlines in Dutch and Irish newspapers about jet aircraft being chartered to fly workers from the west of Ireland to jobs in the Netherlands **because** the Netherlands cannot get workers to do this work.
- NL Onlangs meldden Nederlandse en Ierse kranten dat er vliegtuigen werden gecharterd om arbeiders uit het westen van Ierland naar Nederland te vervoeren voor banen **waar** geen Nederlandse werknemers voor kunnen worden gevonden.

In Hoofdstuk 6 wordt een combinatie van experimentele methodes toegepast om te onderzoeken of beperkende betrekkelijke bijzinsconstructies een interpretatie op discourse-niveau kunnen hebben. De resultaten van twee aanvulexperimenten, een *self-paced reading* experiment en een *eye-tracking* experiment suggereren dat taalgebruikers inderdaad coherentierelaties kunnen infereren tussen beperkende betrekkelijke bijzinnen en hun matrix-zin. Tot de beschikbare interpretaties behoren zowel causale relaties (specifiek: gevolg-oorzaak), zoals in (8), als relaties met een vorm van contrast (specifiek negatief causale verbanden: oorzaak-onverwacht gevolg), zoals in (9).

- (8) De politie is nu op zoek naar de vrouw die de mishandeling moet hebben gezien.⁹
- (9) Jelle moet blokken voor examens die hij al gehaald heeft.¹⁰

Hoewel betrekkelijke beperkende bijzinnen op syntactisch niveau verbonden zijn aan een zelfstandig naamwoord, kunnen ze op discourse-niveau verbonden zijn aan de volledige matrix-zin. Beperkende betrekkelijke bijzinnen zouden daarom niet volledig moeten worden uitgesloten als potentiële discourse-segmenten. Omdat er echter geen aanwijzingen zijn dat alle beperkende betrekkelijke bijzinnen op discourse-niveau aan de volledige matrix-zin verbonden zijn, zal voor nauwkeurige tekstsegmentatie per beperkende betrekkelijke bijzin bepaald moeten worden of deze als zelfstandig segment moet worden aangemerkt. Dit vereist wederom de interpretatie van het fragment.

De linguïstische markering van coherentierelaties

De andere hoofdvraag van dit proefschrift is wanneer een coherentierelatie expliciet gemarkeerd wordt door middel van een connectief. Deze vraag wordt onderzocht in Hoofdstuk 4 en 5. Aan de basis van beide hoofdstukken ligt een geannoteerd parallel

⁹ <http://www.at5.nl/artikelen/182634/kind-met-hoofd-tegen-auto-geslagen-na-grap-met->

¹⁰ <https://consent.hartvannederland.nl/?url=https://www.hartvannederland.nl/nieuws/2018/jelle-moet-blokken-voor-examens-die-hij-al-gehaald-heeft-door-fout-van-school/>

corpus dat bestaat uit ongeveer 2000 Engelse coherentierelaties uit het Europarlorpus (vergaderingen van het Europees Parlement) met vertalingen naar het Nederlands, Duits, Frans en Spaans. Bij het annoteren van deze coherentierelaties is gebruik gemaakt van een annotatiemodel op basis van cognitieve primitieven: CCR (Cognitive approach to Coherence Relations: Sanders, Spooren & Noordman 1992). CCR wordt beschreven, besproken, en uitgebreid in Hoofdstuk 3; het geannoteerde corpus wordt vervolgens gebruikt in Hoofdstuk 4 en 5.

In Hoofdstuk 4 wordt het geannoteerde parallelle corpus gebruikt om de hypothese te toetsen dat taalgebruikers in eerste instantie simpele relaties verwachten en pas een complexere relatie afleiden als de tekst andere indicaties geeft. Cognitief complexe coherentierelaties, zoals de relaties in (1), worden daarom vaker met een connectief gemarkeerd dan simpele coherentierelaties, zoals de relatie in (2). In de studie wordt cognitieve complexiteit bepaald op basis van logica, taalverwerving, taalverwerking en Mental Space Theory. De data zijn grotendeels in lijn met deze hypothese.

In de parallelle corpusstudie in Hoofdstuk 4 wordt tevens rekening gehouden met de positie van een coherentierelatie in de hiërarchische discoursstructuur. In de dataset waren relaties die een andere relatie bevatten of die een segment met een andere relatie delen minder vaak gemarkeerd door een connectief dan relaties die geen andere relatie bevatten of een segment delen. De contrastrelatie tussen de eerste en laatste zin in (3), hieronder herhaald als (10), is een voorbeeld van een relatie die andere relaties bevat (twee causale relaties, gemarkeerd door *namelijk* en *daardoor*). De tweede zin in (10) is een voorbeeld van een segment dat gedeeld wordt door twee relaties; het vormt het tweede segment van de relatie tussen zin 1 en zin 2 en het eerste segment van de relatie tussen zin 2 en zin 3)

- (10) Waterijsjes zijn de beste keus. Hier zit namelijk geen vet in en vaak ook minder suiker dan in chocolade-, room- of yoghurtijs. Daardoor bevatten ze minder calorieën. Ijsjes die minder goed scoren zijn extra grote waterijsjes en ijsjes met room of yoghurt.

Relaties die ingebed zijn in een andere relatie (zoals de relaties gemarkeerd door *namelijk* en *daardoor* in (10) ingebed zijn in een contrastrelatie) of in een syntactische constructie (zoals in (5) de *omdat*-relatie ingebed is in een betrekkelijke bijzin) waren in de dataset vaker expliciet gemarkeerd dan niet-ingebede relaties. Hier was echter sprake van een *confound* met cognitieve complexiteit: complexe relaties waren vaker ingebed dan simpele relaties. In de statistische analyse bleek cognitieve complexiteit een betere voorspeller van explicietheid versus implicietheid. Op basis van Hoofdstuk 4 kan geconcludeerd worden dat of een coherentierelatie met een connectief

gemarkeerd wordt, samenhangt met zowel de cognitieve complexiteit van de relatie als de positie van de relatie in de hiërarchische discourse-structuur.

Connectieven zijn de prototypische markeerders van coherentierelaties, maar taalgebruikers kunnen bij het infereren relaties tussen discourse-segmenten ook gebruikmaken van andere elementen. In Hoofdstuk 5 wordt de bijdrage van andere talige elementen, specifiek elementen binnen de segmenten, aan de markering van coherentierelaties onderzocht. Op basis van een literatuurstudie worden drie manieren geformuleerd waarop connectieven en segment-interne elementen kunnen ‘samenwerken’ in de markering van coherentierelaties: *division of labor*, *agreement* en *general collocation*. Het bestaan van deze drie interacties wordt vervolgens aan een eerste toets onderworpen aan de hand van het geannoteerde parallelle corpus. Bij *division of labor* overlappen het connectief en het segment-interne element in betekenis en maakt de aanwezigheid van de één de aanwezigheid van de ander vaak overbodig. In dit type interactie kan het segment-interne element als signaal voor de coherentierelatie fungeren omdat het de relatie expliciet markeert. Bij *agreement* overlappen het connectief en het segment-interne element ook in betekenis, maar maakt de aanwezigheid van de één de aanwezigheid van de ander niet overbodig (net zoals onderwerp en persoonsvorm allebei getal kunnen uitdrukken). In dit type interactie kan de aanwezigheid van het segment-interne element de verwachting van een coherentierelatie oproepen doordat taalgebruikers dit segment-interne element vaak in dit type coherentierelatie tegenkomen. In *general collocation* is er geen overlap in de betekenis van het connectief en het segment-interne element. Hier kan het segment-interne element als markeerder van coherentierelaties fungeren doordat het vaak in een bepaald type relatie voorkomt en dus de verwachting van die relatie oproept.

De algemene hypothese is dat relaties met een segment-interne markeerder minder vaak een connectief zullen bevatten dan relaties zonder een dergelijk segment-interne element. Segment-interne elementen functioneren echter niet allemaal op dezelfde manier als markeerders van coherentierelaties. Bij *division of labor* is de aanwezigheid van een segment-interne element een veel sterkere voorspeller voor de afwezigheid van het connectief dan bij *agreement* of *general collocation*. In de laatste twee gevallen verhoogt de aanwezigheid van een segment-interne markeerder de kans op de afwezigheid van het connectief, maar dit mechanisme is in deze typen interacties zwakker en waarschijnlijk meer onderhevig aan andere factoren die de expliciete versus impliciete markering van coherentierelaties beïnvloeden.

In Hoofdstuk 4 en 5 worden dus drie mechanismen besproken die lijken te beïnvloeden of een coherentierelatie door middel van een connectief wordt gemarkeerd: de basisverwachting van taalgebruikers om simpele relaties tegen te komen, de positie van de relatie in de hiërarchische discourse-structuur en de aanwezigheid van andere markeerders. Deze factoren sluiten elkaar geenszins uit. Zo zouden sprekers of schrijvers het connectief weg kunnen laten in een cognitief

complexe relatie als een ander talig element de relatie al markeert, of een connectief kunnen toevoegen aan een simpele relatie als deze ingebed is in een andere relatie. Andersom zouden luisteraars of lezers hun basisverwachting van een cognitief simpele relatie kunnen aanpassen als ze een segment-intern element tegenkomen dat een sterke indicator is voor een complexer type coherentierelatie.

Conclusie

Dit proefschrift draagt bij aan onderzoek over discourse-segmentatie en de linguïstische markering van coherentierelaties. Het geeft inzicht in de manier waarop en waarom coherentierelaties expliciet gemarkeerd worden en tussen welke delen van een tekst taalgebruikers coherentierelaties infereren. Daarbij draagt het bij aan de verfijning van richtlijnen voor discourse-segmentatie en -annotatie, beide belangrijke methodologische instrumenten binnen onderzoek naar de coherentie van teksten.

Curriculum Vitae

Jet Hoek was born on July 4, 1990 in Arnhem, The Netherlands. She completed her secondary education at the Stedelijk Gymnasium Arnhem in 2007. She went on to study at the Radboud University in Nijmegen, where she obtained her BA in English Language and Culture in 2010, and her MA in American Studies in 2012 (cum laude). In the spring of 2011, she was a visiting scholar at the University of California at Berkeley. In 2013, she obtained her MA in Linguistics at the Radboud University in Nijmegen (cum laude). She started her PhD project at Utrecht University in April 2014. This dissertation is the result of the research she carried out during that period. While finishing up her PhD project (September 2017-April 2018), Jet also worked part-time as a research assistant on the ACAD project (Automatic Coherence Analysis of Dutch). In October 2018, she will start as a post-doctoral researcher at the University of Edinburgh.