

Cognitive Complexity and the Linguistic Marking of Coherence Relations

A parallel corpus study

Jet Hoek,¹ Sandrine Zufferey,² Jacqueline Evers-Vermeul,¹ & Ted J.M. Sanders¹

¹ Utrecht University, Utrecht Institute of Linguistics OTS

² University of Bern

Abstract

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 paper, 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*.

Keywords

Implicitness, implicitation, expectedness, cognitive complexity, coherence relation, Cognitive approach to Coherence relations (CCR)

Highlights

- The linguistic marking of coherence relations is influenced by cognitive complexity
- The systematicity in the marking of coherence relations is similar across languages
- Uses Cognitive approach to Coherence Relations to determine a relation's complexity
- Both *relational* and *syntactic dependency* appear to be related to *relation type*

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 et al., 1992:2). Such discourse segments consist of minimally a clause.¹

(1) Because [Peter had stolen some cash from his grandma’s purse,] [he was grounded for a month.]

(2) [Jane and Harry did a few jumping jacks] before [they started their morning run.]

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 corpus² (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] [it just stirs up hate.] {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; PDTB Research Group, 2008) and Taboada (2006) on a corpus

¹ See Hoek et al. (2017) 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.

² The ep-number (ep-year-month-day) following each fragment refers to the corpus file from which the fragment was retrieved.

annotated using RST (Rhetorical Structure Theory; Mann & 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 et al., 1996, van Silfhout et al., 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.

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 et al., 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,] but [Kate went home with \$630.]

(5) Although [Jamie had never worked a day in his life,] [he was insanely rich.]

In this paper, 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 et al. (1997a) 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 et al., 1997b). 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 2.4 and 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 paper, 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 2.

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, 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).³

- (6) *I have every confidence that if [we and the politicians of Europe support him] [we will see the kind of reform we need.]* {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] because [they grow crops which have never been subsidised by this Union.]* {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.

1.3 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 2 provides an overview of CCR, discusses the cognitive complexity of the primitives' values, and formulates hypotheses concerning the

³ For a discussion of coherence relations embedded under syntactic constructions and the discourse segment status of syntactically embedded clauses see Hoek, Evers-Vermeul, and Sanders (2017).

marking of coherence relations. Section 3 outlines the parallel corpus study. Results from this study are reported in Section 4, and discussed in Section 5.

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 (Sanders et al, 1992, 1993; Traxler et al., 1997a; Köhne & Demberg, 2013). 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 (cf. Spooren & Sanders, 2008; Evers-Vermeul & Sanders, 2009, 2011; 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 & Sweetser, 2009): expressions that typically establish Mental Spaces. Shifting between mental spaces has been associated with higher costs of processing a coherence relation (Traxler et al., 1997a; Segal & Duchan, 1997; Segal et al., 1991; 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.

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.] Also, [it is important that the core labour standards – these ILO standards – are given the dedication and focus needed for their implementation.]
{ep-02-07-04}
- (9) [I will take the advice of the Commission on this occasion,] but [I hope that we get it right.]
{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 (cf. Clark, 1974; Murray, 1997, Wason & Johnson-Laird, 1972) and acquired later than positive relations (cf. 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) (cf. Pander Maat, 1998; J. Sanders, 1994; Verhagen, 2000, 2005).

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

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.

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 paper, 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,] so [we can vote on it today.] {ep-02-12-17}

(14) [The ECOFIN Council is acting with extreme arrogance] if [it believes it can simply disregard growth and employment in the European Union.] {ep-96-06-18}

(15) [Services such as tourism and clean taxes could register.] In addition, [groups of similar small business could possibly cooperate and register together.] {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 (cf. Bloom et al., 1980; Clark, 2003; Evers-Vermeul & Sanders, 2009; Katz & Brent, 1968; Piaget, 1924/1969), but once acquired, causal relations are processed faster (cf. Haberlandt & Bingham, 1978; Sanders & Noordman, 2000; Mak & Sanders, 2013) 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.⁴ 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 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).

⁴ The *event-indexing model* by Zwaan, Langston, and Graesser (1995) also includes something along the lines of causality being expected by default.

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

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.⁵ In this paper, 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 corpus⁶ (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]
because [it has not been possible to set a safe level of residues due to their toxic effects on
human beings.] {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.

(18) [Workers should be informed at the very start about the situation pertaining to their jobs]
because [they are part and parcel of that company.] {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.] because [it is a
change of views.] {ep-96-04-15}

⁵ The *source of coherence* values were originally called *semantic* and *pragmatic*. These were later renamed as *objective* and *subjective*, respectively.

⁶ This is not entirely surprising, since the corpus consists of transcribed spoken –although in part prepared– discourse.

Epistemic relations tend to be acquired later (cf. Evers-Vermeul & Sanders, 2011; Spooren & Sanders, 2008; van Veen, 2011; Zufferey, 2010) and processed slower (cf. Canestrelli et al. 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 (cf. 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 et al., 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 (cf. 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

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,] [we are going to have to pick up a huge bill.] {ep-03-09-03}

- (22) [It has to be pursued in good times and in bad,] unless [there is a fundamental change in the situation.] {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.

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.

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 implications, 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 implication in translation have also been attributed to structural or linguistic differences between languages in a language pair (Becher, 2011; Fabricius-Hansen, 1999; Hansen-Schirra et al., 2007).

More importantly for the research question addressed in this paper, we assume that regardless of the exact reason for a single instance of implication 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 2008, p.3 and Das & Taboada, 2013, p.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 (Miltasakaki et al., 2004; PDTB Research Group, 2008). 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.

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 et al., 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.⁷ 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 (cf. 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.

⁷ 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/speech act	basic/non-basic
<i>So</i> (197)	positive	causal	content/epistemic/speech act	basic
<i>Also</i> (199)	positive	additive	content/epistemic/speech act	N.A.
<i>In addition</i> (195)	positive	additive	content/epistemic/speech act	N.A.
<i>Although</i> (254)	negative	causal/additive	content/epistemic/speech act	basic/non-basic/N.A.
<i>But</i> (213)	negative	causal/additive	content/epistemic/speech act	basic/non-basic/N.A.
<i>If</i> (254)	positive	conditional	content/epistemic/speech act	basic/non-basic
<i>Unless</i> (216)	negative	conditional	content/epistemic/speech act	basic/non-basic

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. The final dataset contained 1916 original English relations, with translations into Dutch, German, French, and Spanish.

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), which is why the exact annotation of the

relational dependency measures cannot be replicated.⁸ 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.

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 Hoek et al. (2017). Segmentation was done prior to annotation by one annotator (first 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 (first and second author), 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 ($\kappa > 0.7$). One annotator then finished the annotations (first author). In total, 393 (20.5%) relations were double annotated using the CCR primitives.⁹ 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, 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 our primary 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 κ_{MAX} , which is also used to correct for asymmetric marginals, it does not overcorrect in case of

⁸ 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” (Prasad et al. 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.

⁹ 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 3.2.

asymmetric unbalanced marginals, which are mainly caused by annotator biases.¹⁰ 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 (cf. Robaldo & Miltsakaki, 2014; Zufferey & Degand, 2013). 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 (first and second author). 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 (first author). Table 3 shows the overall agreement scores for each dependency measure.

¹⁰ 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

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 implicature 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 implicatures found in our study is much lower than many implicature rates reported in NLP research, compare e.g., Meyer and Webber’s (2013) 18% implicature 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,] [I am happy to go ahead now.] {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,] [we will not be able to make the most of them.] {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] [it is more complex.] {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.,] [it will totally distort the vote on this report.] {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 e.g., 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,] so [the public was able to read in a very clear way where the failings were.] {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,] because [in Ireland at the moment there are 18,500 people employed in this sector.] {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,] [they withdrew the winners of the Coca-Cola Cup, the League Cup, from UEFA's competitions.]

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,] so [there is some contradiction there.] {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] because [we say nastier things or pass nastier resolutions]
 {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 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 2, and the hypotheses of the effects of the different dependency measures, presented in Section 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 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 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 et al., 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*.¹¹ 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.¹² 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%.

¹¹ 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.

¹² 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	p
<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</i>	Δ Negative	1.24	0.25	< .001
<i>Basic operation</i>	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	z-score	p
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 overall number of implicatures varied between target languages, with the least implicatures from English into Spanish; there were fewer implicatures 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 implicature 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.

The model includes an interaction between *polarity* and *basic operation* that is largely in line with the hypothesis formulated in (16) in Section 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 implicated. 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 implicature 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,] [we should encourage international work-sharing to the profit of poorer economies and ourselves.] {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,] [they are meaningless.] {ep-00-03-15}

(34) [We will come back again with material on these issues], because [these problems, in one way or another, are here to stay,] 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.] 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.] {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

5 Discussion

This paper 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) elaboration on 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 paper 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 paper. 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 implicatures in this study (see PDTB Research Group, 2008:65-90), 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 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 paper, 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 et al. (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.

Funding

This work was funded through the SNSF Sinergia project MODERN (CRSII2_147653).

Acknowledgments

We are grateful to Huub van den Bergh for his invaluable statistics advice. In addition, we would like to thank the anonymous reviewers for their insightful comments.

References

- Andersson, Marta & Spenader, Jennifer, 2014. Result and Purpose relations with and without 'so'. *Lingua* 148, 1-27.
- Asr, Fatimeh T. & Demberg, Vera, 2012. Implicitness of discourse relations. *Proceedings of COLING 2012*. Mumbai, India, pp. 2669-2684.
- Barlow, Michael, 2008. Parallel texts and corpus-based contrastive analysis. In: De los Angeles Gómez-González, M. et al. (eds.), *Current Trends in Contrastive Linguistics. Functional and Cognitive Perspectives*. Amsterdam: John Benjamins, pp. 101-121.
- Bates, Elisabeth A., 1976. *Language and Context: The Acquisition of Pragmatics*. New York: Academic Press.
- Becher, Victor, 2011. When and why do translators add connectives? *Target* 23 (1), 26-47.
- 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.
- 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: House, J. & Blum-Kulka, S (Eds.), *Interlingual and Intercultural Communication*. Tübingen: Gunter Narr Verlag, pp. 17-36.
- 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.
- Cartoni, Bruno, Zufferey, Sandrine & Meyer, Thomas, 2013. Using the Europarl corpus for cross-linguistic research. *Belgian Journal of Linguistics* 27 (1), 23-42.
- 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*. The Hague: Mouton. pp. 1291-1498.
- Dancygier, Barbara & Sweetser, Eve E., 2005. *Mental Spaces in Grammar: Conditional Constructions*. Cambridge: Cambridge University Press.
- Das, Debopam, and Taboada, Maite, 2013. Explicit and implicit coherence relations: A corpus study. *Proceedings of the Canadian Linguistic Association (CLA) Conference*. Victoria, Canada. <http://homes.chass.utoronto.ca/~cla-acl/actes2013/actes2013.html>.

- Das, Debopam, Taboada, Maite, and McFetridge, Paul, 2015. RST Signalling Corpus LDC2015T10. Web Download. Philadelphia: Linguistic Data Consortium.
- 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.
- 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, & 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.
- Fabricius-Hansen, Cathrine, 1999. Information packaging and translation: Aspects of sentence splitting (German - English/Norwegian). In: Doherty, M. (Ed.), *Sprachspezifische Aspekte der Informationsverteilung*. Akademie-Verlag, Berlin, pp. 175-214.
- 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.
- Frank, Austin F. & Jaeger, T. Florian, 2008. Speaking rationally: Uniform information density as an optimal strategy for language production. *Proceedings of the 28th meeting of the Cognitive Science Society*. Washington DC, USA, pp. 939-944.
- 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.
- 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.
- Halverson, Sandra, 2004. Connectives as a translation problem. In: Kittel, H. et al. (Eds.), *An International Encyclopedia of Translation Studies*. Berlin/New York: Walter de Gruyter, pp. 562-572.
- 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.
- Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M., 2015. The role of expectedness in the implicitation and explicitation of discourse relations. *Proceedings of the Second Workshop on Discourse in Machine Translation (DiscoMT)*. Lisbon, Portugal, pp. 41-46.

- Hoek, Jet, Evers-Vermeul, Jacqueline, & Sanders, Ted J.M., 2017. Segmenting discourse: Incorporating interpretation into segmentation? *Corpus Linguistics and Linguistic Theory*. Online Preview.
- Hothorn, Torsten, Bretz, Frank, & Westfall, Peter, 2008. Simultaneous Inference in General Parametric Models. *Biometrical Journal* 50 (3), 346-363.
- Katz, Evelyn W. & Brent, Sandor D., 1968. Understanding connectives. *Journal of Verbal Learning and Verbal Behavior* 7, 501-509.
- Kintsch, Walter, 1998. *Comprehension: A Paradigm for Cognition*. Cambridge: Cambridge University Press.
- Koehn, Phillip, 2005. Europarl: A parallel corpus for statistical machine translation. Tenth Machine Translation Summit (MT Summit X), Phuket, Thailand. <http://homepages.inf.ed.ac.uk/pkoehn/publications/europarl-mtsummit05.pdf>.
- Köhne, Judith & Demberg, Vera, 2013. The time-course processing of discourse connectives. *Proceedings of the 35th Annual Meeting of the Cognitive Science Society (CogSci2013)*. Berlin, Germany, pp. 2760-2765.
- Knott, Alistair, 2001. Semantic and pragmatic relations and their intended effects. In: Sanders, T.J.M., Schilperoord, J., & Spooren, W.P.M.S. (Eds.), *Text Representation: Linguistic and Psycholinguistic Aspects*. Amsterdam & Philadelphia: John Benjamins, pp. 127-151.
- 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.
- Kruger, Haidee, 2017. The effects of editorial intervention: implications for studies of the features of translated language. In: De Sutter, G., Lefer, M., and Delaere, I., (Eds.), *Empirical Translation Studies: New Methodological and Theoretical Traditions* (Vol. 300 of TILSM series). The Hague: Mouton de Gruyter, pp. 113-156.
- Levy, Roger & Jaeger, T. Florian, 2007. Speakers optimize information density through syntactic reduction. In: Schlökopf, B., Platt, J., & Hoffman, T. (Eds.), *Advances in Neural Information Processing Systems (NIPS)* (Vol. 19). Cambridge, MA: MIT Press, pp. 849-856.
- 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.
- 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.
- Meyer, Thomas, & Popescu-Belis, Andrei, 2012. Using sense-labeled discourse connectives for statistical machine translation. *Proceedings of the Joint Workshop on Exploiting Synergies*

- between Information Retrieval and Machine Translation (ESIRMT) and Hybrid Approaches to Machine Translation (HyTra). Association for Computational Linguistics. Sofia, Bulgaria, pp. 129-138.
- Meyer, Thomas & Webber, Bonnie L., 2013. Implication of discourse connectives in (machine) translation. Proceedings of the Workshop on Discourse in Machine Translation (DiscoMT). Sofia, Bulgaria, pp. 19-26.
- Miltsakaki, Eleni, Prasad, Rashmi, Joshi, Aravind K, & Webber, Bonnie L., 2004. The Penn Discourse Treebank. Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004). Lisbon, Portugal, pp. 2237-2240.
- 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.
- Noordman Leo G.M. & de Blijzer, Femke, 2000. On the processing of causal relations. In: Couper-Kuhlen, E. & Kortmann, B. (Eds.), *Cause, Condition, Concession and Contrast: Cognitive and Discourse Perspectives*. Berlin/New York: Mouton de Gruyter, pp. 35-56.
- Pander Maat, Henk L.W., 1998. Classifying negative coherence relations on the basis of linguistic evidence. *Journal of Pragmatics* 30 (2), 177-204.
- Patterson, Gary & Kehler, Andrew, 2013. Predicting the presence of discourse connectives. Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing. Seattle, WA., pp. 914-923.
- 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.]
- PDTB Research Group, 2008. The Penn Discourse TreeBank 2.0. Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008). European Language Resources Association (ELRA), pp. 2961-2968.
- Prasad, Rashmi, Miltsakaki, Eleni, Dinesh, Nikhil, Lee, Alan, Joshi, Aravind K., Robaldo, Livio, & Webber, Bonnie L., 2007. The Penn Discourse Treebank 2.0 Annotation Manual. IRCS technical report. http://repository.upenn.edu/cgi/viewcontent.cgi?article=1203&context=ircs_reports (accessed 04.22.2014).
- R Core Team, 2016. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Renkema, Jan, 2004. *Discourse Studies: An Introductory Textbook*. Amsterdam: Benjamins.
- Robaldo, Livio, & Miltsakaki, Eleni, 2014. Corpus-driven semantics of concession: Where do expectations come from? *Dialogue & Discourse*, 5 (1), 1-36.

- Sanders, José, 1994. Perspectives in narrative discourse. PhD thesis, Tilburg University. Available online at <https://pure.uvt.nl/portal/files/896522/305833.pdf>
- Sanders, Ted J.M., 2005. Coherence, causality and cognitive complexity in discourse. *Proceedings/Actes SEM-05, First International Symposium on the Exploration and Modelling of Meaning*. Biarritz, France, pp. 105-114.
- 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: Sanders, T.J.M. & Sweetser, E.E. (Eds.), *Causal Categories in Discourse and Cognition*. Berlin: Mouton de Gruyter, pp. 19-60.
- 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.
- Segal, Erwin M. & Duchan, Judith F., 1997. Interclausal connectives as indicators of structuring in narrative. In: Costermans, J. & Fayol, M. (Eds.), *Processing Interclausal Relationships: Studies in the Production and Comprehension of Text*. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 95-119.
- 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.
- Sperber, Dan & Wilson, Deirdre, 1985. Loose talk. *Proceedings of the Aristotelian Society* LXXXVI, pp. 540-549.
- Spooren, Wilbert P.M.S., 1997. The processing of underspecified coherence relations. *Discourse Processes* 24 (1), 149-168.
- 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.
- 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.

- Trabasso, Tom. & van den Broek, Paul, 1985. Causal thinking and the representation of narrative events. *Journal of Memory and Language* 24 (5), 612-630.
- 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., Ake, Joy P., & Moxey, Linda M., 1997b. Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23 (1), 88-101.
- van den Broek, Paul, 1990. Causal inferences in the comprehension of narrative texts. In: Graesser, A.C. & Bower, G.H. (Eds.), *The Psychology of Learning and Motivation: Inferences and Text Comprehension*. San Diego, CA: Academic Press, pp. 175-194.
- 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: Couper-Kuhlen, E. & Kortmann, B. (Eds.), *Cause, Condition, Concession and Contrast: Cognitive and Discourse Perspectives*. Berlin/New York: Mouton de Gruyter, pp. 361-380.
- Verhagen, Arie, 2005. *Constructions of Intersubjectivity: Discourse, Syntax, and Cognition*. Oxford: Oxford University Press.
- Wason, Peter C. & Johnson-Laird, Philip, 1972. *Psychology of Reasoning: Structure and Content*. Cambridge, MA: Harvard University Press.
- Wilson, Deirde & Sperber, Dan, 2005. Relevance Theory. In: Horn, L. & Ward, G. (Eds.), *The Handbook of Pragmatics*. New York: Wiley, pp. 607-632.
- 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, 2013. Representing the meaning of discourse connectives for multilingual purposes. *Corpus Linguistics and Linguistic Theory*. Online preview.
- Zufferey, Sandrine, & Gygas, 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.