

Implicitness of Discourse Relations

Fatemeh Torabi Asr and Vera Demberg

Cluster of Excellence Multimodal Computing and Interaction (MMCI)

Saarland University

Campus C7.4, 66123 Saarbrücken, Germany

`fatemeh@coli.uni-saarland.de`, `vera@coli.uni-saarland.de`

Abstract

The annotations of explicit and implicit discourse connectives in the Penn Discourse Treebank make it possible to investigate on a large scale how different types of discourse relations are expressed. Assuming an account of the Uniform Information Density hypothesis, we expect that discourse relations should be expressed explicitly with a discourse connector when they are unexpected, but may be implicit when the discourse relation can be anticipated. We investigate whether discourse relations which have been argued to be expected by the comprehender exhibit a higher ratio of implicit connectors. We find support for two hypotheses put forth in previous research which suggest that continuous and causal relations are presupposed by language users when processing consecutive sentences in a text. We then proceed to analyze the effect of Implicit Causality (IC) verbs (which have been argued to raise an expectation for an explanation) as a local cue for an upcoming causal relation.

Keywords: Causality, Continuity, Discourse relations, Discourse cues, Implicit discourse relations, Corpus study, Uniform Information Density.

1 Introduction

David Hume, in his prominent work “An enquiry concerning human understanding” proposed that ideas in the human mind were associated according to at least three types of relations: resemblance, contiguity in time or place, and causality (Hume, 1784). Since then, many language scientists have tried to adapt this idea about human general reasoning to the world of language (Simon, 1952; Hobbs, 1990; Kehler, 2000). A *discourse relation* as defined by linguists is an inference intended by the writer or made by a reader to establish local coherence among individual sentences. For example, a binary *causal* relation refers to a logical cause-consequence inference whose elements are directly accessible in the text: two propositional arguments plus an operator that could be an explicit connective such as “because” (1-a). The operator might sometimes not be explicit in the text, but in most cases a suitable connective could be inserted to specify the coherence relation between the involved propositions (1-b).

- (1) a. Bill took his daughter to the hospital, because she looked pale and sick in the morning.
- b. I was very tired last night. [Therefore] I went to sleep earlier than usual.

According to some cognitive theories on discourse processing, people have expectations about inter-sentential relations when reading a text, which bias their inferential decisions during comprehension (Segal et al., 1991; Murray, 1997; Levinson, 2000; Sanders, 2005). Two important characteristics that are expected to exist between consecutive events in a text are continuity and causality. Segal et al. (1991) and Murray (1997) argue that readers expect a sentence to be causally congruent and continuous with respect to its preceding context. Continuity in the sense of Segal et al. means that the same frame of reference is maintained, for example by subsequent sentences talking about the same event, without shift in perspective. Continuous discourse relations are claimed to be easier to process and more expected than other types. On the other hand, relations that are discontinuous (for example adversatives) would be less expected and more difficult to process. This notion also includes temporal continuity, implying that a non-linearity in presenting a consequence before its effect or any two situations in a reverse temporal order (1-a) is less expected than a relation keeping the forward temporal transition between events (1-b). A second hypothesis about what kind of discourse relations are typically assumed by comprehenders was proposed by Sanders (2005) in his “causality-by-default” hypothesis, which states that language users prefer causal relations to other types (such as mere expansion or temporal relations) when establishing discourse coherence.

To investigate the validity of these hypotheses in the experimental paradigm, researchers have studied participants’ sentence completion preferences, coherence judgments, and reading patterns during reading of the sentence pairs implying different types of discourse relations (Irwin, 1980; Trabasso et al., 1984; Caron et al., 1988; Millis et al., 1995; Murray, 1997; Kuperberg et al., 2011). We are, however, not aware of a large-scale study of these hypotheses in naturally occurring texts. To address this, we analyzed the use of causal/temporal relations and their markers in the Penn Discourse Tree Bank (PDTB; Prasad et al., 2008), a large corpus of newspaper texts which is annotated with discourse relations.

Our methodology relies strongly on the assumption of *implicitness of the discourse connector as a sign of expectation of the discourse relation*: if readers have a default preference to infer a specific relation in the text, this type of relation should tend to appear without explicit markers. The assumption is drawn on a usage-based approach to the study of language preferences that links comprehension phenomena with typicalities in production (Langacker, 2000). It also can be thought of in terms of the Uniform Information Density (UID) hypothesis (Frank and Jaeger, 2008) which suggests that humans tend to spread information evenly across a text or utterance, thereby reducing or omitting redundant optional markers (e.g., Levy and Jaeger, 2007; Florian Jaeger, 2010). At the level of inter-sentential relations, this would mean that the presence of explicit connectives is necessary when the relation is unexpected, but that a connective may be implicit if the relation is predictable. Hence, we investigate the validity of the following predictions in the corpus:

1. continuous discourse relations should be implicit more often than discontinuous ones
2. causal relationships should be implicit more often than other discourse relationships
3. relations which keep the forward temporality should be implicit more often than their backward counterparts.

These predictions refer to general tendencies of comprehenders to expect that phrases and sentences are linked causally, and that the sentences in e.g. a narrative are interpreted in a continuous manner. Beside these general expectations regarding the upcoming discourse relation, expectations can also be influenced by local factors: (Pitler et al., 2008) found that discourse relations in the PDTB are not independently distributed. For example, explicit comparisons are significantly more often followed by implicit contingencies than would be expected under the independence assumption. Furthermore, Sanders (1997) and (Sagi, 2006) showed that the genre of a text, or more specifically, the distribution of different types of relations in a particular text can shape the expectation of a reader about what relations will appear more frequently later in the text.

The discourse relation between two sentences can also become clear by other means such as shared entities, adverbial phrases and even the type of verbs used in the arguments. An interesting case are so-called implicit causality (IC) verbs (as the verb “scolded” in (2-a); see Section 5.6 for more details), which have recently been argued to trigger the expectation for a reason (see (2-b)) to be communicated in the following sentence (Rohde and Horton, 2010), such that the discourse relation between the sentences is a backward causal relation.

- (2) a. Arthur scolded Patricia.
b. She had put thumbtacks on the teacher’s chair.

According to the UID hypothesis, we would expect that causal relations which contain an IC verb in their first argument (Arg1) would be expressed without an explicit connective more often, as the reason relation is predictable. In other words, as our third hypothesis, we investigate whether:

4. backward causal relations that contain an IC verb (which already marks causality) in their first argument are implicit more often than those that contain a non-IC verb.

This work strives to get a better understanding of when discourse relations are expressed implicitly, and when they should be explicitly stated. Furthermore, it can possibly inform work on automatic identification of discourse relations, especially in the absence of explicit discourse markers (Sporleder, 2008; Lin et al., 2009; Zhou et al., 2010). The rest of the paper is structured as follows. Section 2 includes a review of the causality-by-default and the continuity hypotheses and the related experimental studies. In Section 3, we introduce the PDTB corpus and the relevant relations we extract to obtain evidence for the aforementioned hypotheses. Section 4 will describe our method of analysing implicitness and Section 5 includes the results. Finally, conclusions are presented in Section 6.

2 Background on the Continuity and Causality Hypotheses

In this section we explain the *continuity* hypothesis (Segal et al., 1991; Murray, 1997) and the *causality-by-default* hypothesis (Sanders, 2005). Taken together, these hypotheses suggest that language users first try to establish causal relatedness and temporal linearity between phrases when processing a text.

2.1 The Continuity Hypothesis

Levinson (2000) notes in his discussion on presumptive meanings that “*when events are conjoined they tend to be read as temporally successive and if all plausible, as causally linked.*”. An early notion of the continuity hypothesis proposed by Segal et al. (1991) suggests that discourse connectives are used to mark the deictic continuity or discontinuity in texts (their study was limited to narratives). As Murray (1997) states, the hypothesis is that comprehension difficulty ensues when a text event is discontinuous without this discontinuity having been explicitly marked. Segal et al. (1991) found that comprehenders expect subsequent sentences to be causally congruent and temporally continuous. Support for the continuity hypothesis comes also from a series of experiments by Murray (1997). In an initial sentence completion experiment, subjects had to write a continuation to the preceding discourse starting with a connective that was either additive, causal, or adversative. Murray confirmed that sentences generated in response to an additive or causal connective generally depicted continuous events with respect to the preceding discourse, while completions following adversative connectives depicted discontinuous events. In a subsequent experiment, he asked people to read sentence pairs with inappropriately placed connectives, and found that the disruptive effect was largest for inappropriately placed adversative connectives. Sentences with inappropriately placed adversative connectives were also judged as less coherent than ones with inappropriately placed causal or additive connectives. In particular, there was no difference in processing disruption or coherency judgment between inappropriately used additive and causal connectives, which is why Murray attributes the difference between these connectors and adversative ones to the underlying continuity or discontinuity of the described event. Murray interprets his findings in terms of readers generally expecting a continuous event and their processing being more disrupted by a connective that signals an upcoming discontinuity, than by one that would signal continuity. Taken together with an earlier experiment (Murray, 1995), which showed that correctly placed adversative connectors also have a greater beneficiary effect than correctly placed causal or additive connectors, Murray concludes that adversative connectives are more *salient* than connectives that signal continuity.

While he generally classifies causal relations as continuous ones, Murray also notes that a

connective like “because” which signals a temporally non-linear causal relation (*backward* transition from the effect to the cause) should have stronger contextual effects than connectives such as “so” or “therefore”. In the literature, some types of adversative relations are interpreted as *negative causal* relations, e.g., a pair of sentences connected with *although* implies a causal relation in which an unexpected consequence has happened (König, 1991). According to Segal et al. (1991) and Murray (1997), such relations — usually referred to as *concession* — are not expected to the same degree as positive causal relations, which benefit from a higher degree of continuity. The continuity hypothesis also predicts that temporal relations between sentences which cue a non-linear relationship between the arguments would be more difficult to process, and that cues for temporal non-linearity (such as “after”, as opposed to “before”, which indicates the expected temporal order) should be more salient.

2.2 The Causality-by-default Hypothesis

The main motivation behind our study of causal relations comes from Sanders’s cognitive theory of discourse representation. Specifically, the causality-by-default hypothesis states:

“because experienced readers aim at building the most informative representation, they start out assuming the relation between two consecutive sentences is a causal relation” – Sanders (2005) .

Experimental evidence for this claim comes from a range of studies on understanding causal relations in narratives (Trabasso et al., 1984), the effect of connectives on recall of inter-sentential relations (Irwin, 1980; Caron et al., 1988) and preferences in sentence completion tasks (Murray, 1997). A recent study of online sentence processing furthermore reveals that causal relations between sentences facilitate processing even in the absence of discourse connectors: Kuperberg et al. (2011) finds that a small discourse consisting of three sentences was easier to process when the sentences were causally related. Specifically, a larger N400 (an EEG signal which typically indicates semantic anomalies) was found when sentences were irrelevant. All of these findings suggest that readers have a prior expectation that consecutive sentences in a text should be causally related and congruent, unless an explicit cue such as adversative connectives (e.g., *but*) provides marking for another type of relation. The most relevant experiment specific to the tendency towards causal inference is again the one by Murray (1997) in which subjects were asked to continue individual sentences that ended with either a period or a connective of the aforementioned types (additive, causal or adversative). The majority of the answers for the no-connective condition conveyed a causal relation, meaning that subjects often chose a type of continuation that provided a cause or a consequence for the given sentence instead of a simple additive continuation or an adversative one.

However, arguments against the causality-by-default hypothesis can also be found in the literature. Millis et al. (1995) performed an experiment where two consecutive sentences (that did not stand in an obviously causal relationship) were connected with a full stop, or one of the three discourse connectors “because”, “and” or “after”, as the indicators of causal, additive and temporal relations, respectively. The sentence pairs inherently could be interpreted as expressing any of the mentioned relation types. Millis et al. found that causal inferences (as measured by asking participants a “Why?” question after pair of sentences) were only reliably made in the “because” condition, but not in the conditions where the sentences were connected by a period or one of the other connectors. They concluded that

the discourse marker “because” played a very important role in people’s forming of an inference, and that this inference was not formed automatically in these contexts.

The studies we reviewed are all small scale and use carefully designed experimental materials; it is, however, an open question whether the hypotheses generated based on the experimental studies also hold for naturally occurring texts. An additional difficulty is that results from previous studies are often not easily comparable, as they use a slightly different taxonomy of discourse relation types (e.g., *adversative* relations in Murray (1997) includes both negative causal and negative additive relations, while in other related studies such as König (1991), Couper-Kuhlen and Kortmann (2000), and Köhne and Demberg (2011) the former is referred to as a *concession* relation). In this paper, we use the discourse relation categorization of the Penn Discourse Treebank and see how different sentence connectives are being used in naturally occurring text with respect to the extent they reflect causality and continuity.

3 An Overview of the PDTB Corpus

Penn Discourse Tree Bank (PDTB; Prasad et al., 2008) is a large corpus of texts from the Wall Street Journal, which is annotated with discourse relations between every pair of adjacent clauses, sentences or larger text spans¹. The PDTB covers all 25 sections of the Penn Treebank, which has been annotated with various other linguistic information such as syntactic structures and semantic frames.

Discourse relations can either be marked explicitly with a discourse connector (such as “because”, “nevertheless”, “and”), or be implicit². Both explicit and implicit discourse relations are annotated in the corpus, and each discourse relation is marked with a discourse connector and two propositional arguments. In the case of implicit relations, a suitable discourse connector was identified (and inserted) by the annotators. Each argument is an independent text segment whose boundaries are also determined by the annotators. Labeling of the relations has been done according to a hierarchy of discourse relation senses (see Figure 1), including four top-level classes: CONTINGENCY, COMPARISON, TEMPORAL and EXPANSION.

Temporal Relations TEMPORAL relations include *Synchronous* and *Asynchronous* types. *Asynchronous* relations have two subtypes *precedence* and *succession*, which mark forward (3-a) and backward (3-b) temporal transitions respectively.

- (3) a. He believes [that \$55 a share is the most you can pay for Georgia Gulf], before [it becomes a bad acquisition].
- b. [The fields were developed], after [the Australian government decided in 1987 to make the first 30 million barrels from new fields free of excise tax].

We classify *Asynchronous* temporal relations as markers of *discontinuity*, following (Segal et al., 1991). Among the *Asynchronous* temporal relations, the ones where events are in the correct temporal order should be easier to process than the ones where temporal order is

¹For details on the choice of text spans and adjacency we refer the reader to Prasad et al. (2008).

²Other than explicit and implicit relations, the corpus contains two other categories of relations, namely EntRel, indicating the relation between sentences only according to the common entities, as well as AltLex, in which the relation between the two arguments is not lexicalized via the defined set of connectives. We do not include AltLex category in our analysis as it only contains one percent of all the tagged relations.

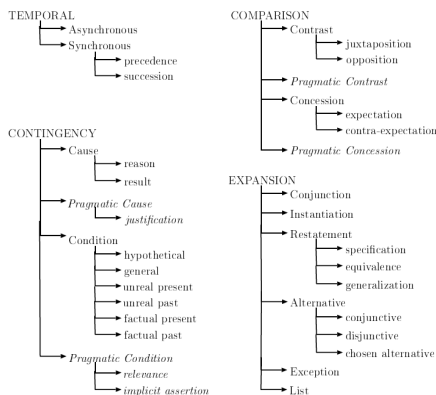


Figure 1: Hierarchy of senses in PDTB (Prasad et al., 2008)

reversed. **Synchronous** temporal relations are harder to classify: they sometimes introduce new events and should thus be classified as discontinuous.

Contingency Relations Causal relations in the PDTB hierarchy are categorized as members of the **CONTINGENCY** family. **Cause** itself is divided into two subtypes, namely **reason** and **result**. These two subtypes, respectively, indicate forward and backward cause-consequence relations between their arguments as shown in (4-b) (**reason**) and (4-a) (**result**).

- (4) a. [The governor couldn't come]_{cause}, so [the lieutenant governor welcomed the special guests]_{consequence}.
 b. [There was some profit-taking]_{consequence}, because [prices for all precious metals had risen to levels at which there was resistance to further advance]_{cause}, he said.
 c. [Mrs Yeargin is lying], because [they found students in an advanced class a year earlier who said she gave them similar help].

In addition to the **Cause** type in the **CONTINGENCY** class, there is a type called **Pragmatic Cause** which is much less frequent but still relevant to our study. It includes only one subtype, namely **justification**, indicating an epistemic causal relation in which the second argument provides a justification for a claim expressed in the first argument, see (4-c). In the literature, this has also been referred to as *diagnostic* relation (Traxler et al., 1997). In PDTB hierarchy conditional discourse relations are also categorized as **CONTINGENCY** relations, see Figure 1. They typically include “if” sentences but are not in the domain of our study, as we cannot classify them as either continuous / discontinuous or causal.

Comparisons All types of **COMPARISON** relations can be classified as discontinuous. As mentioned at the end of Section 2.1, König (1991) proposed that concessive relations are the dual of causal discourse relations in the sense that they involve a negative cause-consequence inference. In the PDTB, **Concession** relations are divided into two subtypes: **expectation** and **contra-expectation**, which represent as the duals of **reason** and **result**

respectively. In **expectation** relations (5-b) the second argument of the connective is a cause for something that is in contrast with the first argument, and in **contra-expectation** (5-a) the first argument of the connective is a cause for something in contrast to what the second argument asserts.

- (5) a. [The demonstrators have been non-violent]_{cause}, but [the result of their trespasses has been to seriously impair the rights of others unconnected with their dispute]_{neg-consequence}.
- b. [Third, oil prices haven't declined]_{neg-consequence}, although [supply has been increasing]_{cause}.

Other relation types of the comparison class, namely **contrast** and **pragmatic contrast** cover some of the adversative relations that have been compared with causal relations in studies such as Murray (1997).

Expansions Among the **EXPANSION** relations we classify as continuous the types **Instantiation**, **Restatement** and **List**. **Alternative** and **Exception** senses are obviously types of discontinuity. The remaining type, namely **Conjunction** relations are used in cases where “Arg2 provides additional, discourse new, information that is not related to Arg1 in any of the ways described for other types of **EXPANSION**.” (The PDTB Research Group, 2008, p. 37). It is difficult to make one classification for the whole class - during manual inspection we found that these relations should sometimes be classified as *discontinuous* because they indicate a deictic shift in entity.

4 Methods

Given the continuity hypothesis, we would expect continuous relations to be very frequent in the corpus, and in particular more frequent than discontinuous relations. Similarly, we expect to see a relatively high frequency of causal discourse relations. An analysis of simple frequencies of total occurrences would however be rather limited, as there are a number of possible confounds. In particular, we would like to be able to test whether there is any evidence for people generating expectations of upcoming discourse relations, as argued by the continuity and causality literature.

To test whether comprehenders in fact do generate expectations about upcoming discourse relations, we draw on the uniform information density (UID) hypothesis (Frank and Jaeger, 2008). The UID hypothesis suggests that humans tend to spread information evenly across a text or utterance, and thereby use linguistic means in order to reduce or omit highly predictable linguistic material (which, because it is predictable, carries only a small amount of information and therefore would lead to a dip in information density if not reduced or omitted). Florian Jaeger (2010) show this effect for the optional “that” in English complement clauses, which they find is omitted more often if the complement clause is predictable from the verb than when it is not. It has been suggested that the function of such a constant rate of information density would be to facilitate information transfer (Genzel and Charniak, 2002). In analyzing the rate of implicit discourse connectors in function of the predictability of the discourse relation according to the continuity and causality-by-default hypotheses, we hypothesize that speakers are also able to dynamically choose to use an explicit discourse connector or drop it in order to achieve UID. In order to measure the ratio of implicit discourse connectors in a discourse relation, we define the

implicitness measure as follows:

$$\text{Implicitness}(\text{relation}) = \frac{\# \text{ of implicit relation}}{\# \text{ relation}}$$

A large value of implicitness (in particular, larger than average implicitness among all relations, which is 0.46) for a particular discourse relation in the corpus would thus indicate that the relation is expressed without the use of a specific discourse connector more often than average. We interpret high values of implicitness as the producer not needing to explicitly specify the relation for the comprehender because it is predictable (or otherwise easily inferable from local cues³). It would also suggest that the annotators of the corpus tended to mark that relation even in the absence of direct textual signals. On the other hand, a small value of implicitness means that the discourse relation is expressed with an explicit discourse cue more often than average, and we would interpret that as the relation being not easily predictable or difficult to process, such that an explicit marker is needed to avoid a peak in information density.

5 Results

This section will first discuss evidence for general patterns of the use of discourse markers, and then proceed to analyze a specific case of IC verbs, which are a local cue of causality.

5.1 Analysis of Global Expectations




Figure 2: Frequency distribution of different types of relations in PDTB corpus: relation types are arranged across the horizontal axis according to their comprising classes: comparison, contingency, expansion and temporal.

Figure 2 depicts the distribution of different relation types over 19,009 explicit and 16,327 implicit considered relations⁴. There are 16 relation types in the hierarchy, some of which

³This notion of *no need for marking* of specific discourse relations such as causal ones reminds us also of Kehler's coherence theory, which indicates that arguments on syntactic decisions should be revisited with respect to the inferences underlying the establishment of discourse relations. For example, Kehler (2000) asserts that in elliptic structures, requirement of parallelism between the two involved discourse segments is important when a *resemblance* relation is targeted. In contrast, when a causal relation is being expressed via an elliptic construction, people understand it with no need of exact structural parallelism to help them with argument identification and alignment.

⁴Not all relations in the PDTB are annotated down to the third level of the hierarchy (this happened

are more frequent than others. The most frequent labels used to annotate the explicit relations (those associated with a textual discourse marker present in the original text) are **Conjunction**, **Contrast**, **Cause** and **Asynchrony**. It is obvious from the chart that the distribution of explicitly marked discourse relations is different from the distribution of implicit occurrence. The most frequent types among implicit relations are in order **Cause**, **Conjunction**, **Restatement**, and **Contrast**. In terms of total frequencies of the different discourse relationships (see Figure 2), **Conjunction**, **Cause** and **Contrast** are the most common relations, and hence could be claimed to be most expected on a pure frequency-based account. However, the relations **Conjunction** and **Contrast** have rather low rates of implicitness (see Figure 3).




Figure 3: Implicitness of different types of relations in PDTB corpus: relation types are arranged across the horizontal axis according to their comprising classes: comparison, contingency, expansion and temporal.

5.2 Evidence for the Continuity Hypothesis

In Section 3 we classified the PDTB discourse relations with respect to continuity or discontinuity of an event in the sense of Segal et al. (1991). In particular, we argued that all of the discourse relations in the **COMPARISON** and **TEMPORAL**⁵ family describe discontinuous events, while **Cause** and **Pragmatic Cause** are continuous. Within the **EXPANSION** family, we argued that **Instantiation**, **Restatement** and **List** are continuous, while **Exception** and possibly **Conjunction** can be viewed as denoting discontinuous events. Figure 3 depicts the implicitness of different level-2 types of discourse relations in the PDTB. We find that the relations that denote continuous events are exactly the ones that have implicitness values larger than average implicitness, while the ones that can be classified as discontinuous are exactly the ones with lower values of implicitness. The PDTB data thus provides strong supporting evidence for the continuity hypothesis. We will get back to more detailed analyzes concerning the concept of temporal continuity in Sections 5.4 and 5.5.

5.3 Evidence for the Causality-by-default Hypothesis

The causality-by-default hypothesis (Sanders, 2005) proposes that people preferentially interpret consecutive sentences as standing in a causal relationship. The data from the

when no consensus on the more detailed classification could be reached among the annotators). In our study, we only included those relations that were annotated down to at least the second level of the hierarchy. Also, some relations are tagged with more than one type. Hence, the sum over occurrences of all types does not necessarily match the total number of argument pairs in the corpus.

⁵Note: **Synchronous** is not straightforward to classify.

PDTB corpus shows that causal relations are not the only ones that are often expressed without an explicit connective, and hence with view on the UID hypothesis are not the only predictable relation (e.g., **Restatement** is a more implicit relation). Nevertheless, causal relations are the most frequent type of implicit discourse relations in the PDTB, see Figure 2. The implicitness of causal discourse relations is significantly higher (at 0.65) than the implicitness of the other frequent discourse relations, in particular **Conjunction** (0.39), **Concession** (0.15), **Asynchronous** (0.25), as well as the average overall relation types (0.46), – all comparisons significant at $p < 0.001$, according to a binomial test. This result supports the hypothesis we constructed on top of the causality-by-default and the UID hypotheses: in the absence of explicit connectives, a causal relation is expected between neighboring sentences.

5.4 Temporal Continuity

Now it is time to compare the subtypes of causal and temporal relations to investigate whether continuity in the temporal ordering of events is implicit or marked explicitly most often. Table 1 includes frequencies of the forward vs. backward causal, concessive and temporal relations in the corpus. As predicted, for each pair of the same type, the forward relation is associated with a higher degree of implicitness. We performed a binomial significance test and obtained $p < 0.001$ for all pairwise comparisons. This result supports the continuity hypothesis in the sense that discourse markers tend to be dropped when the relation between the arguments conform to linearity in time (Murray, 1997).

Type:subtype	Explicit	Implicit	Implicitness	Signif.
Cause:result	752	1704	0.69	} ***
Cause:reason	1488	2467	0.62	
Concession:contra-expectation	804	186	0.19	} ***
Concession:expectation	392	31	0.08	
Asynchrony:precedence	986	536	0.35	} ***
Asynchrony:succession	1101	161	0.12	

***: significant at $p < 0.001$ according to a binomial test

Table 1: Forward and backward occurrences of causal, concessive and temporal relations.

5.5 Textual Order of Arguments

It should be kept in mind that forward/backwardness of the relation (or the connective) in some cases does not correspond with the order in which the arguments of the relation appear in the text. Connectives such as “because” and “although” can take their subordinate clause (Arg2 according to PDTB) to the beginning of a composite sentence, see for example (6).

- (6) Because the drought reduced U.S. stockpiles, they have more than enough storage space for their new crop, and that permits them to wait for prices to rise.

Since all of the occurrences of a particular relation type (e.g., *reason*) in the PDTB are tagged with the same label regardless of the textual order of the arguments, we performed a second analysis. Table 2 presents separate statistics for the Arg1-connective-Arg2 (ordered)

versus the connective-Arg2-Arg1 (reversed) occurrences for the relations **Cause**, **Concession** and **Asynchrony**, which we also focused on in our analysis of temporal continuity. PDTB annotators always put connectives between the arguments of the implicit relations. That is why we only present numbers for the explicit occurrences. Interestingly, there are always more ordinal modifications (the reversed presentation) when a backward relation of any type is being expressed. This implies that even in the presence of the cues, people have a tendency to keep the textual order of the arguments the same as the temporal order in which the associated events happened. In order to conduct a more accurate analysis of the temporal transition, given the information about argument organization we gathered all implicit and explicit occurrences of temporally linear vs. non-linear instances from all the 6 relation subtypes and performed a correlation analysis. In this analysis reversed occurrences of backward relations count as temporally linear and reversed occurrences of forward relations are taken as temporally non-linear (e.g., a reason relation in which Arg2 appears first in the text is taken as linear, just like a result relation). The chi-square test of the temporal linearity and the implicitness of the relation shows a significant correlation between these two factors ($\chi^2 = 67.31$, $df = 1$, $p < 0.001$). This is a more accurate result (compared with our analysis of continuity merely according to the relation types in 5.4) which indeed supports our hypothesis based on the UID and the continuity hypotheses: temporal forwardness is implicit in relations between consecutive sentences and its explicit cues are typically dropped.

Type:subtype (explicit only)	Ordered	Reversed	Signif.
Cause:result	746	6	} ***
Cause:reason	1324	164	
Concession:contra-expectation	791	13	} ***
Concession:expectation	183	209	
Asynchrony:precedence	931	55	} ***
Asynchrony:succesion	867	234	

***: significant at $p < 0.001$ according to a binomial test

Table 2: Distribution of textually ordered vs. reversed occurrences of arguments in causal, concessive and temporal relations with explicit connectors.

5.6 A potential local predictor: Implicit Causality Verbs

Implicit Causality verbs (such as *adore*, *inspire*, *humiliate*) have been studied for many years, mostly in the context of coreference (whether the subject of the clause explaining the reason is the subject or object of the IC verb). Recently, IC verbs have however also been argued to make comprehenders anticipate an upcoming causal discourse relation, i.e. comprehenders expect to learn about the reason if they hear a sentence with an IC verb, like “Peter adored his older brother”: In a visual world experiment, Rohde and Horton (2010) compared sentences with IC verbs to sentences with transfer-of-possession (TOP) verbs, and found that participants are much more likely to expect a reason following an IC verb sentence than when they have heard a sentence with a TOP verb. This indicates that people are able to take into account local cues like IC verbs and anticipate upcoming discourse relations.

To test on the corpus data the validity of the hypothesis that IC verbs lead people to anticipate causal relationships, we extracted all sentences from the PDTB which contained

an IC verb in the Arg1 of a discourse relation. In order to identify the IC verbs, we used a list of 300 IC verbs provided by Ferstl et al. (2011). In order to avoid noise in the data, we extracted only those instances where the IC verb was the main verb of an Arg1 which contained only a single sentence. Therefore, we simultaneously queried the PDTB annotation and the syntactic annotation of the Penn Treebank. To make sure that the IC verb worked as a cue in the sentence, we only considered relations with ordered arguments, namely the Arg1-connective-Arg2 occurrences. We found that the discourse relation was labeled as **reason** significantly more often if Arg1 contained an IC verb than when it did not ($p < 0.01$), however, the size of the effect was small: likelihood of **reason** given an IC verb in Arg1 was 14.0%, and 11.7% for other verbs. This is support (though relatively weak) for IC verbs actually affecting the upcoming discourse relation.

The more interesting question in the context of the UID hypothesis, however, is whether markers for causal relations following IC verbs are more likely to be absent, due to the added predictability of the **reason** relationship. We compared the implicitness of reason relations where the Arg1 contains an IC verb to the implicitness of reason relations with non-IC verbs as the head of Arg1. Counterintuitively, we found that the implicitness of reason relations with an IC verb in the Arg1 was smaller than for non-IC Arg1s. To make sure that this effect was not due to noise, we also checked all occurrences of IC verbs manually and only included the correct verbs and not homonyms (such as for *lie*) or other unintended semantic sense (e.g. “leave it up to somebody” instead of “leave somebody”). We found that such incorrectly tagged verbs were evenly distributed among implicit and explicit relations, such that the implicitness value was not affected by the noise (implicitness was 61% for IC verb reason relations and 65% for non-IC verb reason relations). Table 3 shows the manually checked numbers within reason relations (for all relations, only automatically extracted numbers are available due to the large number of occurrences, and absence of effect from manually checking the IC verbs in reason relations⁶).

	Total	IC verb in Arg1
Implicit: reason relations	2462	153 (manually checked)
Explicit: reason relations	1324	96 (manually checked)
Implicit: all relations	15682	910 (automatically extracted)
Explicit: all relations	16147	1034 (automatically extracted)

Table 3: Total frequency of relations and the frequency of IC verbs appearing as the head of a single-sentence Arg1. All selected relations are also filtered to have ordered arguments.

The result goes against our predictions and merits further investigation in further work: all verbs need to be classified into their IC semantic class according to the previous research, and finer-grained verb sense disambiguation should also be considered⁷, and further factors such as whether the verb occurs in passive voice could be taken into account. Also, the list

⁶The number of automatically extracted reason relations including IC verbs were 164 and 108 for the implicit and explicit occurrences, respectively.

⁷Ferstl et al. (2011) categorize IC verb usages into 4 classes: AgentPatient, AgentEvocator, Experiencer-Stimulus, and StimulusExperiencer. For example, the AgentPatient class covers *activity* transitive verbs such as “*carry*” which associate Agent and Patient roles to the involved entities. The 300 IC verb list that we employed also did not contain information about which fine-grained verb sense should be treated as an IC verb.

of 300 IC verbs from (Ferstl et al., 2011) are only annotated for their subject or object bias. In future work, these verbs should all be tested for the strength of predicting an upcoming reason relationship; and the result could be taken into account to see whether the prediction of higher implicitness of discourse cues signaling a reason relationship following IC verbs possibly holds for those IC verbs which strongly predict a reason relationship.

6 Conclusions

We conducted an empirical study of discourse relations in newspaper text, specifically the PDTB treebank, with respect to the causality-by-default and continuity hypotheses. We found supporting evidence for both hypotheses: As the *continuity hypothesis* (Murray, 1997; Segal et al., 1991) predicts, discourse relations which are discontinuous or temporally non-linear are much more likely to be expressed with an explicit discourse marker than those which are continuous. The statistics on the forward vs. backward temporal transition between arguments of discourse relations furthermore show a higher degree of implicitness for the forward directionality of all causal, concessive and temporal relations than for the backward versions of them.

Causal relations constitute the largest proportion of the implicit discourse relations in the corpus, which suggests that they are more expected when no discourse marker is present, compared to many other relation types such as temporals, adversatives and additives. From a usage-based perspective, this provides partial support for the *causality-by-default hypothesis* put forth by Sanders (2005) in the sense that causal relations are identified even if no textual element explicitly marks them. However, in the absence of explicit sentence connectives other types of discourse relations could also be inferred, such as restatement or instantiation, which also account for a large proportion of unmarked relations in the PDTB corpus.

We also investigated implicit causality verbs which have been argued to act as local cues for an upcoming causal relationship. However, we found that their presence in the first sentence of a reason relationship does not increase the probability of the explicit connective to be dropped. This finding stands in contrast to what we predicted via an account of the UID hypothesis, which suggests that optional markers might be dropped if they contribute less information. Nevertheless, we observed that presence of an IC verb in the first argument of a sentence pair could generally signal a reason relation in the corpus, in line with the experimental finding of (Rohde and Horton, 2010) on discourse-level predictions. Taken together, our findings raise an interesting question for the future work: to what extent can global patterns vs. local cues account for the discourse relations being left implicit?

While the patterns we observed in the production data are compatible with the mentioned hypotheses about causality and continuity, they do not give us insight about the source of the tendency. Our results along with the related experimental findings can be considered from a frequency-based perspective, meaning that typical patterns in language production lead to expectations during comprehension about causality and continuity. Alternatively, it could be that people have an intrinsic tendency towards congruent and temporally ordered relations both in production and interpretation.

References

- Caron, J., Micko, H. C., and Thuring, M. (1988). Conjunctions and the recall of composite sentences. *Journal of Memory and Language*, 27(3):309–323.

- Couper-Kuhlen, E. and Kortmann, B. (2000). *Cause, Condition, Concession, Contrast: Cognitive and Discourse Perspectives*, volume 33. Walter de Gruyter.
- Forstl, E., Garnham, A., and Manouilidou, C. (2011). Implicit causality bias in english: a corpus of 300 verbs. *Behavior Research Methods*, 43(1):124–135.
- Florian Jaeger, T. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61(1):23–62.
- Frank, A. and Jaeger, T. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. *Proceedings of the 28th meeting of the Cognitive Science Society*.
- Genzel, D. and Charniak, E. (2002). Entropy rate constancy in text. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 199–206.
- Hobbs, J. (1990). *Literature and Cognition*, volume 21. Center for the Study of Language and Information (CSLI lecture notes).
- Hume, D. (1784). *An Enquiry Concerning Human Understanding*. New York: The Liberal Arts Press, 1955 edition.
- Irwin, J. (1980). The effects of explicitness and clause order on the comprehension of reversible causal relationships. *Reading Research Quarterly*, pages 477–488.
- Kehler, A. (2000). Coherence and the resolution of ellipsis. *Linguistics and Philosophy*, 23(6):533–575.
- Köhne, J. and Demberg, V. (2011). Incremental and predictive discourse based on causal and concessive discourse markers – a visual study. In *24th Annual CUNY Conference on Human Sentence Processing*.
- König, E. (1991). Concessive relations as the dual of causal relations. *Semantic Universals and Universal Semantics. Groningen-Amsterdam Studies in Semantics*, 12:190–209.
- Kuperberg, G., Paczynski, M., and Ditman, T. (2011). Establishing causal coherence across sentences: An ERP study. *Journal of Cognitive Neuroscience*, 23(5):1230–1246.
- Langacker, R. (2000). A dynamic usage-based model. *Usage-based Models of Language*, pages 1–63.
- Levinson, S. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. The MIT Press.
- Levy, R. and Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In *Advances in Neural Information Processing Systems*.
- Lin, Z., Kan, M., and Ng, H. (2009). Recognizing implicit discourse relations in the penn discourse treebank. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 343–351.
- Millis, K., Golding, J., and Barker, G. (1995). Causal connectives increase inference generation. *Discourse Processes*, 20(1):29–49.

- Murray, J. (1995). Logical connectives and local coherence. *Sources of Coherence in Reading*, pages 107–125.
- Murray, J. (1997). Connectives and narrative text: The role of continuity. *Memory and Cognition*, 25(2):227–236.
- Pitler, E., Raghupathy, M., Mehta, H., Nenkova, A., Lee, A., and Joshi, A. (2008). Easily identifiable discourse relations. *Technical Reports (CIS)*, page 884.
- Prasad, R., Dinesh, N., Lee, A., Miltsakaki, E., Robaldo, L., Joshi, A., and Webber, B. (2008). The Penn Discourse Treebank 2.0. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008)*, pages 2961–2968.
- Rohde, H. and Horton, W. (2010). Why or what next? eye movements reveal expectations about discourse direction. In *Proceedings of 23rd Annual CUNY Conference on Human Sentence Processing*, pages 18–20.
- Sagi, E. (2006). Context and the processing of discourse: Priming and genre effects on discourse comprehension. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*.
- Sanders, T. (1997). Semantic and pragmatic sources of coherence: On the categorization of coherence relations in context. *Discourse Processes*, 24(1):119–147.
- Sanders, T. (2005). Coherence, causality and cognitive complexity in discourse. In *Proceedings/Actes SEM-05, First International Symposium on the Exploration and Modelling of Meaning*, pages 105–114.
- Segal, E., Duchan, J., and Scott, P. (1991). The role of interclausal connectives in narrative structuring: Evidence from adults’ interpretations of simple stories. *Discourse Processes*, 14(1):27–54.
- Simon, H. (1952). On the definition of the causal relation. *The Journal of Philosophy*, 49(16):517–528.
- Sporleder, C. (2008). Lexical models to identify unmarked discourse relations: Does WordNet help? *Lexical-Semantic Resources in Automated Discourse Analysis*, page 20.
- The PDTB Research Group (2008). The Penn Discourse Treebank 2.0 annotation manual. Technical report, Institute for Research in Cognitive Science, University of Pennsylvania.
- Trabasso, T., Secco, T., and van den Broek, P. (1984). Causal cohesion and story coherence. *Learning and Comprehension of Text*, pages 83–111.
- Traxler, M., Sanford, A., Aked, J., and Moxey, L. (1997). Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(1):88.
- Zhou, Z., Lan, M., Niu, Z., Xu, Y., and Su, J. (2010). The effects of discourse connectives prediction on implicit discourse relation recognition. In *Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 139–146.