
Attribution and its annotation in the Penn Discourse TreeBank

Rashmi Prasad* — Nikhil Dinesh* — Alan Lee* — Aravind Joshi*
— Bonnie Webber**

* *Institute for Research in Cognitive Science (IRCS)*
3401 Walnut Street, Philadelphia, PA 19104

rjprasad@seas.upenn.edu

nikhild@seas.upenn.edu

aleewk@seas.upenn.edu

joshi@seas.upenn.edu

** *University of Edinburgh*
Edinburgh, EH8 9LW Scotland

bonnie@inf.ed.ac.uk

ABSTRACT. In this paper, we describe an annotation scheme for the attribution of abstract objects (propositions, facts, and eventualities) associated with discourse relations and their arguments annotated in the Penn Discourse TreeBank. The scheme aims to capture both the source and degrees of factuality of the abstract objects through the annotation of text spans signalling the attribution, and of features recording the source, type, scopal polarity, and determinacy of attribution.

RÉSUMÉ. Dans cet article, nous décrivons un schéma d'annotation pour l'encodage des objets abstraits (propositions, faits et possibilités) associés aux relations de discours et à leurs arguments tels qu'annotés dans le Penn Discourse TreeBank. Ce schéma a pour objet la capture de la source et du degré de factualité des objets abstraits. Les aspects clés de ce schéma comprennent l'annotation des intervalles textuels signalant l'attribution, ainsi que l'annotation des propriétés caractérisant la source, le type, la polarité de la portée, et le degré de détermination de l'attribution.

KEYWORDS: Attribution, discourse relations, information extraction, Penn Discourse Treebank.

MOTS-CLÉS : Attribution, relations de discours, extraction d'information, Penn Discourse Treebank.

1. Introduction

News articles typically contain a mixture of information presented from different perspectives, and often in complex ways. Writers may present information as known to them, or from someone else's perspective, while further distinguishing between, for example, whether that perspective involves an assertion or a belief. Recent work has shown the importance of recognizing such perspectivization of information for several NLP applications, such as information extraction, summarization, question answering (Wiebe *et al.*, 2004; Stoyanov *et al.*, 2005; Riloff *et al.*, 2005) and generation (Prasad *et al.*, 2005). This allows such applications to distinguish between factual and non-factual information, and to identify the source of the information. Annotation schemes (Wiebe *et al.*, 2005; Wilson and Wiebe, 2005; PDTB-Group, 2006) encode such distinctions to facilitate accurate recognition and representation of such perspectivization of information.

This paper describes an annotation scheme for marking the attribution of discourse relations and their arguments annotated in the Penn Discourse TreeBank (PDTB) (Miltsakaki *et al.*, 2004; Prasad *et al.*, 2004; Webber *et al.*, 2005), the primary goal being to capture the source and degrees of factuality of *Abstract Objects* (AOs) (Asher, 1993). Since discourse relations and their arguments have AO interpretations, several scenarios can be distinguished: in some cases, a discourse relation and its arguments may be attributed to the writer or some other agent introduced in the text; in other cases, the discourse relation is established by the writer, with one or both arguments attributed to others. Applications concerned with recognizing and representing AOs would therefore benefit from understanding how and to whom the components of discourse relations are attributed. Our annotation scheme captures four salient properties of attribution:

- (a) *Source*, which distinguishes between agents to whom AOs are attributed,
- (b) *Type*, which reflects the degree of factuality of the AO,
- (c) *Scopal polarity* of attribution, which indicates polarity reversals of attributed AOs due to surface negated attributions, and
- (d) *Determinacy* of attribution, which indicates the presence of contexts canceling the entailment of attribution.

The scheme also annotates the *text spans* that convey the source of the attribution. The described scheme is an extension of the core scheme used for annotating attribution in the first release of the PDTB (Dinesh *et al.*, 2005; PDTB-Group, 2006). Section 2 gives an overview of the PDTB, Section 3 presents the extended annotation scheme for attribution, Section 4 describes the annotation methodology and a preliminary evaluation, Section 5 discusses some potential uses of the annotation for applications, and Section 6 presents the summary.

2. The Penn Discourse TreeBank (PDTB)

The PDTB contains annotations of discourse relations and their arguments on the Wall Street Journal (WSJ) corpus (Marcus *et al.*, 1993). Following the approach towards discourse structure in Webber *et al.* (2003), the PDTB takes a lexicalized approach towards the annotation of discourse relations, treating *discourse connectives* as the anchors of the relations and thus as discourse-level predicates taking two *Abstract Objects* as their arguments. For example, in (1), *since* is a discourse connective that denotes a TEMPORAL relation between the event of the earthquake hitting and a state where no music is played by a certain woman.¹

(1) *She hasn't played any music since the earthquake hit.* (0766)

There are primarily two types of connectives in the PDTB: “explicit” and “implicit”, the latter being *inserted* between adjacent paragraph-internal sentence pairs not related by an explicit connective.

2.1. Explicit connectives

Explicit connectives refer to lexical items used to trigger discourse relations, and are identified from primarily four grammatical classes:

- (a) Subordinating conjunctions, both bare (e.g., *because*, *when*) and modified (e.g., *only because*, *particularly since*);
- (b) Subordinators (e.g., *in order that*);
- (c) Coordinating conjunctions (e.g., *and*, *or*);
- (d) Discourse adverbials (e.g., *however*, *otherwise*).^{2,3}

Annotation of explicit connectives consists of identifying and recording the text span corresponding to the connective. In Example (1), for instance, *since* is identified as an explicit connective, and annotated by selection of the corresponding text span. In the examples in this paper, explicit connectives are underlined.

2.2. Arguments of connectives

Annotation of the arguments of connectives consists of identifying and recording the text spans that provide the interpretation of the AO arguments of the connective. Because there are, as yet, no generally accepted abstract semantic categories for classifying the arguments to discourse connectives, similar to those for verbs (e.g., *agent*,

1. The 4-digit number in parentheses at the end of an example is its WSJ file number.

2. *Discourse adverbials* are distinguished from *clausal adverbials* (Forbes-Riley *et al.*, 2006).

3. Discourse markers such as *well*, *anyway*, *now*, etc., that signal the organizational structure of the discourse, are not annotated.

patient, theme, etc.), arguments of connectives are simply labelled Arg2, for the argument appearing in the clause syntactically bound to the connective, and Arg1, for the other argument. In the examples here, the text whose interpretation is the basis for Arg1 appears in *italics*, while that of Arg2 appears in **bold**. Example (1) shows that the main clause of the sentence was annotated as Arg1 of the explicit connective *since*, and the subordinate clause was annotated as Arg2. Note that Arg1 and Arg2 can appear in either order, or one may appear interposed with the other. For the subordinating conjunctions, since the subordinate clause is bound to the connective, Arg2 corresponds to the subordinate clause, and hence the linear order of the arguments can be Arg1-Arg2 (Example 1), Arg2-Arg1 (Example 2), or Arg2 may appear embedded in Arg1 (Example 3), depending on the relative position of the subordinate clause with respect to its governing matrix clause.

- (2) Michelle lives in a hotel room, *and* although **she drives a canary-colored Porsche**, *she hasn't time to clean or repair it.* (2402)
- (3) *Most oil companies*, **when they set exploration and production budgets for this year**, *forecast revenue of \$15 for each barrel of crude produced.*⁴ (0725)

The order of the arguments for adverbials and coordinating conjunctions is typically Arg1-Arg2 since Arg1 usually appears in the prior discourse. But as Example (4) shows, the arguments of discourse adverbials can appear embedded within one another. In this example, Arg1 is embedded in Arg2.

- (4) As an indicator of the tight grain supply situation in the U.S., market analysts said **that late Tuesday the Chinese government**, *which often buys U.S. grains in quantity*, **turned instead to Britain to buy 500,000 metric tons of wheat.** (0155)

The basic unit for the realization of an AO argument of a connective is the clause, tensed or untensed, but it can also be associated with multiple clauses, within or across sentences. *Nominalizations* and *discourse deictics* (*this, that*), which can also be interpreted as AOs, can serve as the argument of a connective too. To constrain the amount of text selected for arguments, a *minimality principle* requires an argument to contain the minimal amount of information needed to complete the interpretation of the relation. Any other span of text that is perceived to be relevant (but not necessary) to the interpretation is annotated as *supplementary information*, labelled Sup1, for material supplementary to Arg1, and Sup2, for material supplementary to Arg2.

4. As this example shows, annotations in the PDTB can be discontinuous. Discontinuous annotation is possible for connectives as well, such as for *on the one hand . . . on the other hand*.

2.3. Implicit connectives

Not all discourse relations are triggered by lexical items: In the PDTB, adjacency is taken to be a non-lexical trigger for discourse relations, where the adjacent elements are sentences unrelated by an explicit connective. In such cases, annotators attempt to *infer* a discourse relation, *inserting* as annotation those “implicit” connectives that *best* convey the inferred relations. For example, in (5), annotators infer that the second sentence is related to the first via a CAUSAL relation, (i.e. Mr. Breeden’s wise perception of the ways of Washington is being used to explain the assertion that he may be able to succeed). This inferred CAUSAL relation has been annotated with *because* as the implicit connective.⁵ Implicit connectives together with their sense classification are shown here in small caps.⁶

- (5) *Also unlike Mr. Ruder, Mr. Breeden appears to be in a position to get somewhere with his agenda.* Implicit = BECAUSE (CAUSE) **As a former White House aide who worked closely with Congress, he is savvy in the ways of Washington.** (0955)

Cases where a suitable implicit connective could not be annotated between adjacent sentences are annotated as either:

- (a) **EntRel**, where the second sentence only serves to provide some further description of an entity in the first sentence (Example 6);
- (b) **NoRel**, where no discourse relation or entity-based relation can be inferred; or
- (c) **AltLex**, where the insertion of an implicit connective leads to *redundancy*, due to the discourse relation being *alternatively lexicalized* by some “non-connective” expression (Example 7). Such an expression (shown here in square brackets) anchors the AltLex annotation, similar to explicit connectives.

- (6) *C.B. Rogers Jr. was named chief executive officer of this business information concern.* Implicit = EntRel **Mr. Rogers succeeds J.V. White, 64, who will remain chairman and chairman of the executive committee.** (0929)
- (7) *One in 1981 raised to \$2,000 a year from \$1,500 the amount a person could put, tax-deductible, into the tax-deferred accounts and widened coverage to people under employer retirement plans.* Implicit = AltLex (CONSEQUENCE) **[This caused] an explosion of IRA promotions by brokers, banks, mutual funds and others.** (0933)

5. Implicit connectives between adjacent sentences across paragraphs, as well as relations occurring intra-sententially (as with *free adjuncts*) are currently not annotated.

6. For the first release of the PDTB (PDTB-Group, 2006), implicit connectives have been classified broadly into seven semantic types (CAUSE, CONSEQUENCE, CONDITION, ADDITIONAL-INFO, TEMPORAL, RESTATEMENT/SUMMARIZATION, and CONTRAST). A more fine-grained classification will be followed for the final release. Semantic classification of explicit connectives is also planned (Miltsakaki *et al.*, 2005) for the final release.

2.4. Summary of annotations in PDTB-1.0.

The first release of the Penn Discourse TreeBank, PDTB-1.0 (PDTB-Group, 2006), is freely available from <http://www.seas.upenn.edu/~pdtb>.

PDTB-1.0 contains 100 distinct types of explicit connectives, with a total of 18,505 tokens, annotated across the entire WSJ corpus (25 sections). Implicit connectives have been annotated in three sections (Sections 08, 09, and 10) for the first release, totalling 2,003 tokens. The corpus also includes a broadly defined sense classification for the implicit connectives (see Fn. 6), and attribution annotation with the earlier core scheme. Subsequent releases of the PDTB will include implicit connectives annotated across the entire corpus, attribution annotation using the extended scheme described in this paper, and fine-grained sense classification for both explicit and implicit connectives.

3. Annotation of attribution

Recent work (Wiebe *et al.*, 2005; Prasad *et al.*, 2005; Riloff *et al.*, 2005; Stoyanov *et al.*, 2005) has shown the importance of recognizing and representing the source and factuality of information in certain NLP applications. Information extraction systems, for example, would perform better by prioritizing the presentation of factual information, and multi-perspective question answering systems would benefit from presenting information from different perspectives.

Most of the annotation approaches tackling these issues, however, are aimed at performing classifications at either the document level (Pang *et al.*, 2002; Turney, 2002), or the sentence or word level (Wiebe *et al.*, 2004; Yu and Hatzivassiloglou, 2003). In contrast to these approaches, the focus here is on marking attribution on more analytic semantic units, namely the *Abstract Objects* (AOs) associated with predicate-argument discourse relations annotated in the PDTB, with the aim of providing a compositional classification of their factuality. The scheme isolates four key properties of attribution, to be annotated as features:

- (a) *Source*, which distinguishes between different types of agents (Section 3.1);
- (b) *Type*, which encodes the nature of the relationship between agents and AOs, thereby reflecting their factuality (Section 3.2);
- (c) *Scopal polarity*, which is marked when surface negated attribution reverses the polarity of the attributed AO (Section 3.3) and
- (d) *Determinacy*, which signals a context that cancels what would otherwise be an entailment of attribution (Section 3.4).

In addition, to further facilitate the task of identifying attribution, the scheme also annotates the *text span* signaling attribution (Section 3.5), with the goal of highlighting the textual anchors of the features mentioned above.

Results from preliminary annotations show that a significant proportion (34%) of the annotated discourse relations – for both explicit and implicit connectives – have some non-writer agent as the source, for either the relation or one or both arguments. Thus one cannot simply attribute discourse relations and their arguments to the writer of news text, without being wrong two-thirds of the time. The annotations also show that there are a variety of configurations in which the components of the relations are attributed to different sources, suggesting that recognition of attributions may be a complex task for which an annotated corpus may be useful. For example, in some cases, a relation together with its arguments is attributed to the writer or some other agent, whereas in other cases, while the relation is attributed to the writer, one or both of its arguments is attributed to different agent(s). We are hopeful that the PDTB annotations can provide a useful gold standard of judgments and features, for use in automatically classifying the attribution source.

3.1. Source

The *source* feature distinguishes between:

- (a) the writer of the text (“Wr”),
- (b) some specific agent introduced in the text (“Ot” for other), and
- (c) some arbitrary (“Arb”) individual(s) indicated via a non-specific reference in the text.

In addition, since attribution can have scope over an entire relation, arguments can be annotated with a fourth value “Inh”, to indicate that their source value is inherited from the relation.

Given this scheme for source, there are broadly two possibilities. In the first case, a relation and both its arguments are attributed to the same source, either the writer, as in (8), or some other agent (here, Bill Biedermann), as in (9).⁷

- (8) Since the British auto maker became a takeover target last month, its ADRs have jumped about 78%. (0048)

	REL	Arg1	Arg2
[Source]	Wr	Inh	Inh

- (9) “The public is buying the market when in reality there is plenty of grain to be shipped,” said Bill Biedermann, Allendale Inc. director. (0192)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh

7. Attribution feature values assigned to examples are shown below each example; REL stands for discourse relation; attribution text spans are shown boxed.

As Example (8) shows, text spans for implicit writer attributions (corresponding to implicit communicative acts such as “*I write*”, or “*I say*”) are not marked and are taken to imply writer attribution by default (see also Section 3.5).⁸

In the second case, one or both arguments have a different source from the relation. In (10), for example, the relation and Arg2 are attributed to the writer, whereas Arg1 is attributed to another agent (here, Mr. Green). On the other hand, in (11) and (12), the relation and Arg1 are attributed to the writer, whereas Arg2 is attributed to another agent (the purchasing agents in (11) and Mr. Guterman in (12)).

- (10) When Mr. Green won a \$240,000 verdict in a land condemnation case against the State in June 1983, he says *Judge O’Kicki unexpectedly awarded him an additional \$100,000.* (0267)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Inh

- (11) *Factory orders and construction outlays were largely flat in December* while purchasing agents said **manufacturing shrank further in October.** (0178)

	REL	Arg1	Arg2
[Source]	Wr	Inh	Ot

- (12) *There, on one of his first shopping trips, Mr. Paul picked up several paintings at stunning prices. He paid \$2.2 million, for instance, for a still life by Jan Jansz. den Uyl that was expected to fetch perhaps \$700,000. The price paid was a record for the artist. (...) Afterward, Mr. Paul is said by Mr. Guterman **to have phoned Mr. Guterman, the New York developer selling the collection, and gloated.** (2113)*

	REL	Arg1	Arg2
[Source]	Wr	Inh	Ot

Example (13) shows an example of a non-specific “Arb” source indicated by an agentless passivized attribution on Arg2 of the relation. Note that passivized attributions can also be associated with a specific source when the agent is explicit, as shown in (12), where the explicit agent is Mr. Guterman.⁹ “Arb” sources are also identified

8. It is also possible for an “Ot” attribution to be implicit for a relation or argument. These, however, are inferred from some explicit occurrence of the source in the prior text, and their attribution spans are marked extra-sententially (see Section 3.5).

9. In passivized attributions (e.g., in Examples (12) and (13)), the subject of the infinitive raised to the position of main clause subject is included in the attribution text span. This is because of the convention of including in the attribution span all non-clausal complements and modifiers of the attribution predicate (Section 3.5).

by the occurrences of adverbs like *reportedly*, *allegedly*, etc. Example (14) illustrates one such case.

- (13) Although index arbitrage is said **to add liquidity to markets**,
John Bachmann, . . . says *too much liquidity isn't a good thing.* (0742)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Arb

- (14) *East Germans rallied as officials* reportedly **sought Honecker's ouster.**
 (2278)

	REL	Arg1	Arg2
[Source]	Wr	Inh	Arb

We conclude this section by noting that “Ot” is used to refer to any specific individual as the source. No further annotation is provided to indicate who the “Ot” agent in the text is. Furthermore, as shown in Examples (15-16), multiple “Ot” sources within the same relation do not indicate whether or not they refer to the same or different agents. However, we assume that the text span annotations for attribution, together with an independent mechanism for named entity recognition and anaphora resolution can be effectively exploited to identify and disambiguate the appropriate references.

- (15) *Suppression of the book*, Judge Oakes observed, *would operate as a prior restraint and thus involve the First Amendment.* Moreover, and here Judge Oakes went to the heart of the question, **"Responsible biographers and historians constantly use primary sources, letters, diaries, and memoranda."** (0944)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Ot

- (16) *The judge was considered imperious, abrasive and ambitious*,
those who practiced before him say *...Yet, despite the judge's imperial bearing, no one ever had reason to suspect possible wrongdoing,*
says John Bognato, president of Cambria County's bar association. (0267)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Ot

3.2. Type

The *type* feature signifies the nature of the relation between an agent and an AO, leading to different inferences about the degree of factuality of the AO. We start by making the well-known distinction of AOs into four sub-types: *assertion propositions*, *belief propositions*, *facts* and *eventualities*.¹⁰ This initial distinction is significant in our scheme since it corresponds, in part, to the types of attribution relations and the verbs that convey them, and simultaneously allows for a semantic, compositional approach to the annotation and recognition of factuality.¹¹

3.2.1. Assertion proposition AOs and belief propositions AOs

Proposition AOs involve attribution to an agent of his/her commitment towards the truth of a proposition. A further distinction captures differences in the degree of that commitment, by distinguishing between “assertions” and “beliefs”.

Assertion proposition AOs are associated with a communication type of attribution (“Comm” for short), conveyed by standard verbs of communication (Levin, 1993) such as *say*, *mention*, *claim*, *argue*, *explain* etc. In Example (17), the attribution on Arg1 takes the value “Comm” for *type*. Implicit writer attributions, as with the relation in Example (17), also take the (default) value “Comm”. Note that when an argument’s attribution source is not inherited (as for Arg1 in this example) it takes its own independent value for *type*. This example thus conveys that there are two different attributions expressed within the discourse relation, one for the relation and the other for one of its arguments, and that both involve propositional assertions.

- (17) When Mr. Green won a \$240,000 verdict in a land condemnation case against the State in June 1983, he says *Judge O’Kicki unexpectedly awarded him an additional \$100,000.* (0267)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Inh
[Type]	Comm	Comm	Null

In the absence of an independent occurrence of attribution on an argument, as for Arg2 of Example (17), a “Null” value for the *type* on the argument means that it needs to be derived by independent (here, undefined) considerations under the scope of the relation. Note that unlike the “Inh” value of the *source* feature, “Null” does not indicate inheritance. In a subordinate clause, for example, while the relation denoted by the subordinating conjunction may be asserted, the clause content itself may be presupposed, as seems to be the case for the relation and Arg2 of (17). However, we

10. This corresponds roughly to the top-level tier in the AO hierarchy of Asher (1993).

11. Note that discourse relations are also taken to denote a special class of propositions, called *relational propositions* (Mann and Thompson, 1988) and are themselves treated as abstract objects in the PDTB (Prasad *et al.*, 2005).

found these differences difficult to determine at times, and consequently leave this undefined in the current scheme.

Belief proposition AOs are associated with a "belief" type of attribution, conveyed by propositional attitude verbs (Hintikka, 1971) such as *believe*, *think*, *expect*, *suppose*, *imagine*, etc. This type of attribution is thus called "PAtt" for short. An example of a belief attribution is given in (18).

- (18) Mr. Marcus believes *spot steel prices will continue to fall through early 1990 and then reverse themselves.* (0336)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	PAtt	Null	Null

3.2.2. Fact AOs

Facts AOs involve attribution to an agent of an evaluation towards or knowledge of a proposition whose truth is taken for granted (i.e., a presupposed proposition). Fact AOs are associated with a "factive" type of attribution ("Ftv" for short), conveyed by "factive" and "semi-factive verbs" (Kiparsky and Kiparsky, 1971; Karttunen, 1971) such as *regret*, *forget*, *remember*, *know*, *see*, *hear*, etc. An example of a factive attribution is given in (19). In the current scheme, this class does not distinguish between the true factives and semi-factives, the former involving an attitude/evaluation towards a fact, and the latter involving knowledge of a fact.

- (19) The other side, he argues knows *Giuliani has always been pro-choice, even though he has personal reservations.* (0041)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	Ftv	Null	Null

3.2.3. Eventuality AOs

While **eventuality AOs** are independent of attribution, when they occur with attribution it conveys an agent's intention/attitude towards a considered event, state or action. Eventuality AOs occur with "control" types of attribution ("Ctrl" for short), conveyed by any of three different classes of control verbs (Sag and Pollard, 1991). The first kind is anchored by a *verb of influence* like *persuade*, *permit*, *order*, and involve one agent influencing another agent to perform (or not perform) an action. The second kind is anchored by a *verb of commitment* like *promise*, *agree*, *try*, *intend*, *refuse*, *decline*, and involve an agent committing to perform (or not perform) an action. The third kind is anchored by a *verb of orientation* like *want*, *expect*, *wish*, *yearn*, and involve desire, expectation, or some similar mental orientation towards some state(s) of affairs. These sub-distinctions are not encoded in the annotation, but we have used

the definitions as a guide for identifying these predicates. Note that the syntactic term *control* is used here because these verbs denote uniform structural control properties, but the primary basis for their definition is nevertheless semantic. An example of the control attribution relation anchored by a verb of influence is given in (20).¹²

- (20) Eward and Whittington had planned to leave the bank earlier, but Mr. Craven had persuaded them to remain until **the bank was in a healthy position.** (1949)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	Ctrl	Null	Null

3.3. *Scopal polarity*

The *scopal polarity* feature is annotated on relations and their arguments to identify cases when verbs of attribution are negated on the surface - syntactically (e.g., *didn't say*, *don't think*) or lexically (e.g., *denied*), but when the negation in fact **reverses** the polarity of the attributed relation or argument content (Horn, 1978). Example (21) illustrates such a case. The *but* clause entails an interpretation such as “I think it's not a main consideration”, for which the negation must take narrow scope over the embedded clause rather than the higher clause. In particular, the interpretation of the CONTRAST relation denoted by *but* requires that Arg2 should be interpreted under the scope of negation.

- (21) “*Having the dividend increases is a supportive element in the market outlook,* but I don't think **it's a main consideration,**” he says. (0090)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Ot
[Type]	Comm	Null	PAtt
[Polarity]	Null	Null	Neg

To capture such entailments with surface negations on attribution verbs, an argument of a connective is marked “Neg” for *scopal polarity* when the interpretation of the connective requires the surface negation to take semantic scope over the lower argument. Thus, in Example (21), *scopal polarity* is marked as “Neg” for Arg2. When the neg-lowered interpretations are not present, *scopal polarity* is marked as the default “Null” (such as for the relation and Arg1 of Example 21).

12. While our use of the term *source* applies literally to agents responsible for the truth of a proposition, we continue to use the same term for the agents for facts and eventualities. Thus, for facts, the *source* represents the bearers of attitudes/knowledge, and for considered eventualities, the *source* represents the bearer of intentions/attitudes.

Note that this surface negation can be interpreted as taking scope only over the relation, rather than any argument as well. Since we have not observed this yet in the PDTB, we describe this case with the constructed example in (22). What the example shows is that in addition to entailing (22b) – in which case it would be annotated parallel to Example (21) above – (22a) can also entail (22c), such that the negation is interpreted as taking semantic scope over the relation (Lasnik, 1975), rather than one of the arguments. As the *scopal polarity* annotations for (22c) show, lowering of the surface negation to the relation is marked as “Neg” for the *scopal polarity* of the relation.

- (22) a. John doesn't think *Mary will get cured because she took the medication.*
- b. John thinks *that because Mary took the medication, she will not get cured.*

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	PAtt	Null	Null
[Polarity]	Null	Neg	Null

- c. John thinks *that Mary will get cured not because she took the medication (but because she has started practising yoga.)*

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	PAtt	Null	Null
[Polarity]	Neg	Null	Null

We note that *scopal polarity* does not capture the appearance of (opaque) internal negation that may appear on arguments or relations themselves. For example, a modified connective such as *not because* does not take “Neg” as the value for *scopal polarity*, but rather “Null”. This is consistent with our goal of marking *scopal polarity* only for lowered negation, i.e., when surface negation from the attribution is lowered to either the relation or argument for interpretation.

3.4. Determinacy

The *determinacy* feature captures the fact that the attribution over a relation or argument can **itself** be cancelled in particular contexts, such as within the scope of negation or a conditional. Such indeterminacy is indicated by the value “Indet”, while determinate contexts are simply marked by the default “Null”. The annotation in Example (23) illustrates a case of indeterminacy of the (belief) attribution on the relation.

Here, it is not that a belief or opinion about “our teachers educating our children better if only they got a few thousand dollars a year more” is being attributed to anyone, even “Arb” (ie, an arbitrary individual). Rather, the attribution is only being conjectured as a possibility. This indeterminacy is created by the infinitival context in which the attribution is embedded.

- (23) It is silly libel on our teachers to think *they would educate our children better if only they got a few thousand dollars a year more.* (1286)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	PAtt	Null	Null
[Polarity]	Null	Null	Null
[Determinacy]	Indet	Null	Null

3.5. Attribution spans

In addition to annotating the properties of attribution in terms of the features discussed above, we also annotate the *text span* associated with the attribution. The text span is annotated as a single (possibly discontinuous) complex reflecting the annotated features. The attribution span also includes all non-clausal modifiers of the elements contained in the span, for example, adverbs and appositive NPs. Connectives, however, are excluded from the span, even though they function as modifiers. Example (24) shows a discontinuous annotation of the attribution, where the parenthetical *he argues* is excluded from the attribution phrase *the other side knows*, corresponding to the factive attribution.

- (24) The other side, he argues, knows *Giuliani has always been pro-choice, even though he has personal reservations.* (0041)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	Ftv	Null	Null
[Polarity]	Null	Null	Null
[Determinacy]	Null	Null	Null

We note that in annotating the attribution span as a single complex, we assume that the text anchors of the individual elements of the attribution - the *source*, *type*, *scopal polarity* and *determinacy* - can be identified by independent means from the semantic role annotations (namely, Propbank (Kingsbury and Palmer, 2002)) on the Penn Treebank.

Spans for implicit writer attributions are left unmarked since there is no corresponding text that can be selected. The absence of a span annotation is simply taken to reflect writer attribution, together with the “Wr” value on the source feature.

Recognizing attributions is not trivial since they are often left unexpressed in the sentence in which the AO is realized, and have to be inferred from the prior discourse. For example, in (25), the relation and its arguments in the third sentence are attributed to Larry Shapiro, but this attribution is implicit and must be inferred from the first sentence. The spans for such implicit “Ot” attributions mark the text that provides the inference of the implicit attribution, which is just the closest occurrence of the explicit attribution phrase in the prior text.

- (25) “There are certain cult wines that can command these higher prices,”
says Larry Shapiro of Marty’s, . . . “What’s different is that it is happening
 with young wines just coming out. *We’re seeing it partly because **older
 vintages are growing more scarce.***” (0071)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	Comm	Inh	Null
[Polarity]	Null	Null	Null
[Determinacy]	Null	Null	Null

The final aspect of the span annotation is that we also annotate non-clausal phrases as the anchors of attribution, such as prepositional phrases like *according to X*, and adverbs like *reportedly*, *allegedly*, *supposedly*. One such example is shown in (26). Note that while a specific individual is identified as the source of Arg1 in this example, with “Ot” as the source value, many such phrases, especially the adverbs, refer to a non-specific generic source. In the latter case, the source value is marked as “Arb”. Also, note that the type and scopal polarity of the attribution indicated by such phrasal attributions are assumed to be provided by the phrase itself. In (26), the *according to* preposition head of the attribution phrase is taken to reflect an assertion by the indicated agent, and the type is thus marked as “Comm”.

For phrasal attributions, since the PDTB argument annotation guidelines do not allow for non-clausal modifiers of an argument to be excluded from the selection – a convention – they also appear as part of the argument span they modify. While this is slightly awkward, it does not conflict with any aspect of the PDTB annotation.

- (26) *No foreign companies bid on the Hiroshima project,*
according to the bureau. **But the Japanese practice of deep discounting
 often is cited by Americans as a classic barrier to entry in Japan’s
 market.** (0501)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Inh
[Type]	Comm	Comm	Null
[Polarity]	Null	Null	Null
[Determinacy]	Null	Null	Null

3.6. Attribution of implicit relations

Implicit connectives and their arguments in the PDTB are also marked for attribution. Implicit connectives express relations that are inferred by the reader. In such cases, the writer intends for the reader to infer a discourse relation. As with explicit connectives, implicit relations intended by the writer of the article are distinguished from those intended by some other agent introduced by the writer. For example, while the implicit connective in Example (27) is attributed to the writer, in Example (28), both Arg1 and Arg2 have been expressed by someone else whose speech is being quoted: in this case, the implicit connective is attributed to the other agent.

- (27) *The gruff financier recently started socializing in upper-class circles.* Implicit = FOR EXAMPLE (ADD.INFO) *Although he says he wasn't keen on going, **last year he attended a New York gala where his daughter made her debut.*** (0800)

	REL	Arg1	Arg2
[Source]	Wr	Inh	Inh
[Type]	Comm	Null	Null
[Polarity]	Null	Null	Null
[Determinacy]	Null	Null	Null

- (28) *"We asked police to investigate why they are allowed to distribute the flag in this way. Implicit = BECAUSE (CAUSE) **It should be considered against the law,**"* said Danny Leish, a spokesman for the association. (0814)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	Comm	Null	Null
[Polarity]	Null	Null	Null
[Determinacy]	Null	Null	Null

Attribution is also annotated for AltLex relations but not for EntRel and NoRel, since the former but not the latter indicate the presense of discourse relations.

4. Annotation procedure and evaluation

To date, annotation with the scheme presented in this paper has been carried out by a single annotator on approximately one-half of the explicit connectives annotated in PDTB-1.0 - on 9,000 tokens. The annotator was provided with the four features of attribution along with the set of their corresponding possible features values, to be marked for each discourse relation and each argument. That is, for each discourse relation and its arguments, the task of the annotator was to select from the given set of

feature values a single value for each of the four features of *source*, *type*, *polarity*, and *determinacy*. Training for the annotation was provided both in terms of the definitions in Section 3 as well as in terms of illustrative examples.

When the annotation of attribution is complete for all explicit and implicit connectives, we will evaluate its reliability in terms of inter-annotator agreement. For this, we intend to use a random sampling method. We have prototyped this method on a small, randomly drawn sample of 100 explicit connectives. Agreement was computed for the token as a whole so that a single disagreement for any feature value or attribution span was counted as disagreement for the complete token. With this simple measure, we achieved 84% agreement. The sample drawn for this evaluation is obviously not statistically significant, but it does seem suggestive of what we can expect from a more large-scale evaluation when the annotation is completed. In the full evaluation, we intend to follow a more fine-grained metric where agreement will be computed as a composite of each feature value for the relation as well as its arguments.

5. Applications: future work

As we mentioned at the outset, certain applications, for example information extraction, question answering and generation, are concerned with categorizing information units in terms of their source and factuality. Considering Abstract Objects as the information units as we have for the annotation of attribution here, the two primary questions such applications would like answered are whether or not the AO information is factual or non-factual, and what its source is.

The goal of the attribution annotation carried out in the PDTB is to provide a source for learning to recognize the contexts which lead to such inferences. One of the first tasks towards identifying attribution contexts is to recognize the attribution spans that scope over the abstract objects. In the simplest case, such recognition distinguishes between writers and non-writers as the *sources* since writer attributions are most often implicit. Note that in addition to differentiating between the types of sources, such cases can also convey different degrees of factuality, since writer attributions would typically be taken to convey factual information, much more so than non-writer attributions. Furthermore, the latter admits to further possible differences in the degree of factuality through the *type* feature. For example, speech attributions (expressed with “verbs of communication”) have a higher degree of factuality than belief attributions (expressed with “verbs of propositional attitude”), while factive attributions (expressed with “factive verbs”) are most likely associated with the same degree of factuality as writer attributions, if not more. Our hypothesis is that all of these distinctions are learnable from the attribution spans (or the absence thereof) since they contain linguistic expressions of these inferences.

The *determinacy* feature is also an indicator of non-factuality. Systems that extract information units below the level of the sentence, in particular clausal units denoting Abstract Objects would be amiss if they labelled the information contained in the

conditional relation in Example (23) – repeated here as (29) – as factual. What is crucial about this example is *not* that the belief attribution expressed with the verb *think* assigns the information a lower degree of factuality (than, say, a writer attribution, a factive attribution, or a speech attribution), but rather that the attribution itself, whatever *type* it may be, is cancelled by virtue of being embedded in (in this case) an infinitival context. Such inferences of indeterminacy allow for the conclusion that the information is hypothetical at best.

- (29) It is silly libel on our teachers to think *they would educate our children better if only they got a few thousand dollars a year more.* (1286)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Inh
[Type]	PAtt	Null	Null
[Polarity]	Null	Null	Null
[Determinacy]	Indet	Null	Null

Another significant use of the annotation of attribution is towards identifying the arguments of relations. In many cases, an attribution appearing over an argument of a relation has only an evidential role (Dinesh *et al.*, 2005), as in Example (17) – repeated here as (30) – where the attribution on Arg1 is outside the scope of the temporal relation. Applications concerned with identifying the argument structure of discourse relations need to be able to recognize the evidential role of such attributions and exclude them while extracting the argument structures.

- (30) When Mr. Green won a \$240,000 verdict in a land condemnation case against the State in June 1983, he says *Judge O’Kicki unexpectedly awarded him an additional \$100,000.* (0267)

	REL	Arg1	Arg2
[Source]	Wr	Ot	Inh
[Type]	Comm	Comm	Null

Simple exclusion of the attribution span in such cases is not always sufficient, however, since certain elements appearing within the attribution span on the surface seem to be necessary for determining the semantic argument of the relation. This occurs with surface negated attribution in particular, which in our annotation scheme is captured with the *scopal polarity* feature. For instance, in Example (21) – repeated here as (31) – what is important to recognize for extracting the arguments of the CONTRAST relation denoted by *but* is that while the Arg2 attribution lies outside the scope of the relation, the negation appearing with the associated attribution phrase must be retained to take scope over Arg2. Without the retention of this negation, the necessary textual entailment for interpreting the relation is lost. Learning to recognize such contexts and determining the correct scope of negation for such entailments is thus a necessary task for argument identification algorithms.

- (31) “*Having the dividend increases is a supportive element in the market outlook, but I don’t think it’s a main consideration,*” he says. (0090)

	REL	Arg1	Arg2
[Source]	Ot	Inh	Ot
[Type]	Comm	Null	PAtt
[Polarity]	Null	Null	Neg

6. Summary

In this paper, we have proposed and described an annotation scheme for marking the attribution of explicit and implicit discourse connectives and their arguments in the Penn Discourse TreeBank. We defined the notion of attribution as relations between individuals and Abstract Objects, and presented the scheme in detail with examples, outlining the “feature-based annotation” in terms of the *source*, *type*, *scopal polarity*, and *determinacy* associated with attribution, and the “span annotation” to highlight the text reflecting the attribution features. We described the annotation procedure and a small evaluation experiment for determining inter-annotator agreement. Finally, we discussed the use of the annotations for natural language applications concerned with identifying the source and factuality of information units, and also discussed the role of the *polarity* annotations for determining correct textual entailments in certain cases of surface negated attributions.

Acknowledgements

The Penn Discourse TreeBank project is partially supported by NSF Grant: Research Resources, EIA 02-24417 to the University of Pennsylvania (PI: A. Joshi). We are grateful to the anonymous reviewers for their useful comments.

7. References

- Asher N., *Reference to Abstract Objects in Discourse*, Kluwer, Dordrecht, 1993.
- Dinesh N., Lee A., Miltsakaki E., Prasad R., Joshi A., Webber B., “Attribution and the (Non)-Alignment of Syntactic and Discourse Arguments of Connectives”, *Proceedings of the ACL Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*, Ann Arbor, Michigan, 2005.
- Forbes-Riley K., Webber B., Joshi A., “Computing Discourse Semantics: The Predicate-Argument Semantics of Discourse Connectives in D-LTAG”, *Journal of Semantics*, vol. 23, p. 55-106, 2006.
- Hintikka J., “Semantics for propositional attitudes”, in L. Linsky (ed.), *Reference and Modality*, Oxford, p. 145-167, 1971.
- Horn L., “Remarks on Neg-raising”, in P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*, Academic Press, New York, 1978.

- Karttunen L., "Some observations on factivity", *Papers in Linguistics*, vol. 4, p. 55-69, 1971.
- Kingsbury P., Palmer M., "From TreeBank to PropBank", *Proceedings of Third International Conference on Language Resources and Evaluation (LREC-2002)*, Las Palmas, Canary Islands, Spain, 2002.
- Kiparsky C., Kiparsky P., "Fact", in D. D. Steinberg, L. A. Jakobovits (eds), *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics and Psychology*, Cambridge University Press, Cambridge, p. 345-369, 1971.
- Lasnik H., "On the semantics of negation", *Contemporary Research in Philosophical Logic and Linguistic Semantics*, Dordrecht: D. Reidel, p. 279-313, 1975.
- Levin B., *English Verb Classes And Alternations: A Preliminary Investigation*, University of Chicago Press, 1993.
- Mann W. C., Thompson S. A., "Rhetorical structure theory. Toward a functional theory of text organization", *Text*, vol. 8, n° 3, p. 243-281, 1988.
- Marcus M. P., Santorini B., Marcinkiewicz M. A., "Building a Large Annotated Corpus of English: The Penn Treebank", *Computational Linguistics*, vol. 19, n° 2, p. 313-330, 1993.
- Miltsakaki E., Dinesh N., Prasad R., Joshi A., Webber B., "Experiments on Sense Annotation and Sense Disambiguation of Discourse Connectives", *Proceedings of the Fourth Workshop on Treebanks and Linguistic Theories (TLT-2005)*, Barcelona, Spain, 2005.
- Miltsakaki E., Prasad R., Joshi A., Webber B., "Annotating discourse connectives and their arguments", *Proceedings of the HLT/NAACL Workshop on Frontiers in Corpus Annotation*, Boston, MA, 2004.
- Pang B., Lee L., Vaithyanathan S., "Thumbs up? Sentiment classification using machine learning techniques", *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2002)*, Philadelphia, PA, 2002.
- Prasad R., Joshi A., Dinesh N., Lee A., Miltsakaki E., Webber B., "The Penn Discourse TreeBank as a resource for Natural Language Generation", *Proceedings of the Corpus Linguistics Workshop on Using Corpora for NLG (UCNLG-2005)*, Birmingham, U.K., 2005.
- Prasad R., Miltsakaki E., Joshi A., Webber B., "Annotation and Data Mining of the Penn Discourse TreeBank", *Proceedings of the ACL Workshop on Discourse Annotation*, Barcelona, Spain, 2004.
- Riloff E., Wiebe J., Phillips W., "Exploiting subjectivity classification to improve information extraction", *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-2005)*, Pittsburgh, PA, 2005.
- Sag I. A., Pollard C., "An Integrated Theory of Complement Control", *Language*, vol. 67, n° 1, p. 63-113, 1991.
- PDTB-Group, The Penn Discourse TreeBank 1.0 Annotation Manual, Technical Report n° IRCS-06-01, Institute for Research in Cognitive Science, University of Pennsylvania, 2006.
- Stoyanov V., Cardie C., Wiebe J., "Multi-Perspective Question Answering Using the OpQA Corpus", *Proceedings of the Joint Conference of Human Language Technology and Empirical Methods in Natural Language Processing (HLT-EMNLP-2005)*, Vancouver, Canada, 2005.
- Turney P. D., "Thumbs up or Thumbs down? Semantic orientation applied to unsupervised classification of reviews", *Proceedings of the Association for Computational Linguistics (ACL-2002)*, Philadelphia, PA, 2002.

- Webber B., Joshi A., Miltsakaki E., Prasad R., Dinesh N., Lee A., Forbes K., “A Short Introduction to the PDTB”, *Copenhagen Working Papers in Language and Speech Processing*, 2005.
- Webber B., Joshi A., Stone M., Knott A., “Anaphora and Discourse Structure”, *Computational Linguistics*, vol. 29, n° 4, p. 545-587, 2003.
- Wiebe J., Wilson T., Bruce R., Bell M., Martin M., “Learning subjective language”, *Computational Linguistics*, vol. 30, n° 3, p. 277-308, 2004.
- Wiebe J., Wilson T., Cardie C., “Annotating expressions of opinions and emotions in language”, *Language Resources and Evaluation*, vol. 39, n° 2-3, p. 165-210, 2005.
- Wilson T., Wiebe J., “Annotating attributions and private states”, *Proceedings of the ACL Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*, Ann Arbor, Michigan, 2005.
- Yu H., Hatzivassiloglou V., “Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences”, *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003)*, Sapporo, Japan, 2003.