

GRaSP: A Multilayered Annotation Scheme for Perspectives

Chantal van Son, Tommaso Caselli, Antske Fokkens, Isa Maks,
Rosier Morante, Lora Aroyo and Piek Vossen

Vrije Universiteit Amsterdam

De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

{c.m.van.son, t.caselli, antske.fokkens, isa.maks, r.morantevallejo, lora.aroyo, piek.vossen}@vu.nl

Abstract

This paper presents a framework and methodology for the annotation of perspectives in text. In the last decade, different aspects of linguistic encoding of perspectives have been targeted as separated phenomena through different annotation initiatives. We propose an annotation scheme that integrates these different phenomena. We use a multilayered annotation approach, splitting the annotation of different aspects of perspectives into small subsequent subtasks in order to reduce the complexity of the task and to better monitor interactions between layers. Currently, we have included four layers of perspective annotation: events, attribution, factuality and opinion. The annotations are integrated in a formal model called GRaSP, which provides the means to represent instances (e.g. events, entities) and propositions in the (real or assumed) world in relation to their mentions in text. Then, the relation between the source and target of a perspective is characterized by means of perspective annotations. This enables us to place alternative perspectives on the same entity, event or proposition next to each other.

Keywords: source perspective, multilayered annotation, attribution

1. Introduction

With the Internet having secured an increasingly prominent place in society, the current age is characterized by the ability of individuals to transfer and access information about the world. The textual form in which this information is usually represented is rich and complex. Texts contain information about what is happening in the world, where, when, and who is involved. At the same time, they are a reflection of ongoing debates in our society, stances on particular issues (e.g. abortion, vaccinations, etc.), and interpretative frames on events and their causes (e.g. conspiracy theories on 9/11). Textual data always provide specific perspectives of the author and quoted sources on the information they contain. Mining information from texts thus implies dealing with these perspectives.

In the last decade, different aspects of linguistic encoding of perspectives have been targeted as separated phenomena through different annotation initiatives, each with its own approaches and goals. Targeted aspects of perspectives include, for example, attribution (Prasad et al., 2007; Pareti, 2012), factuality (Saurí and Pustejovsky, 2009; Diab et al., 2009), and opinion (Wiebe et al., 2005). Coordinating initiatives such as the Unified Linguistic Annotation project¹ have tried to technically combine such annotations into a unique annotation model, but they lack an overarching framework for the various layers of annotation from different resources. Furthermore, annotation initiatives such as those proposed by Prasad et al. (2007) and Pareti (2012) have attempted to tackle the annotation of perspectives in a unified approach, but with different levels of success.²

In our approach, the notion of perspectives lies at the semantic-pragmatic interface. A perspective of a given source consists of a statement that is made (what is said or

believed, and what not), the factuality attributed to a source and the opinion or sentiment the source expresses towards the target (if any). This paper presents a framework that addresses the most problematic issues with respect to the annotation and alignment of aspects of perspectives.

Our main contributions are the following. First, the annotation scheme we propose is multilayered, i.e. the annotations are split into small subsequent subtasks. We expect that keeping the tasks small, clear and simple will lead to higher inter-annotator agreement. In addition, this allows us to better monitor interactions between different semantic and pragmatic phenomena involved in perspectives. Second, the annotations are integrated into a formal model that provides the means to represent instances (e.g. events, entities) and propositions in the (real or assumed) world in relation to their mentions in text. Then, the relation between the source and target of a perspective is characterized by means of perspective-related annotations such as attribution, factuality and opinion. This enables us to place alternative perspectives on the same object next to each other, as well as to investigate world views of sources through their statements. Third, we address the exact scope of perspective values. For instance, the default interpretation of *Harry was not killed with a knife* is that Harry was killed, but not with a knife. Our model allows to represent this scope of negation by assigning different factuality values to the killing of Harry on the one hand, and the killing being carried out with a knife on the other hand.

The remainder of this paper is structured as follows. First, we review an overview of previous annotation efforts aimed at capturing different aspects of perspectives in Section 2. The fundamental notions and approach underlying the proposed annotation scheme are introduced in Section 3, which explains the main elements of perspectives, the formal model in which our annotations are integrated, and our motivation for a multilayered approach. Section 4 describes the four layers that we have currently defined for the an-

¹<https://catalog.ldc.upenn.edu/LDC2009T07>

²For example, Kohen's kappa on attribution type and scopal change is 0.64 and 0.61 respectively, while it reaches 0.87 on the identification of attribution relations (Pareti, 2012).

notation of perspectives: events, attribution, factuality, and opinion. Finally, we conclude and summarize our work and give an outlook on future work in Section 5.

2. Related Work

In our annotations, events play an important role because we consider them to be the basic semantic elements that may give rise to or be involved in perspectives. A well-known specification language for events is TimeML (Pustejovsky et al., 2003a), which has been consolidated as an international cross-language ISO standard (Pustejovsky et al., 2010) and has been used as the annotation language for the TempEval shared task series (Verhagen et al., 2009). Its reference corpus is TimeBank (Pustejovsky et al., 2003b). TimeML defines an event as something that can be said to obtain or hold true, to happen or to occur. TimeML adopts a surface-based annotation of texts and morpho-syntactic information plays a key role for detecting all possible mentions of an event. According to TimeML, both *demonstrations* and *taken (place)* in Example 1 are to be annotated as valid event mentions.

1. Several pro-Iraq **demonstrations** have **taken** place in the last week.

In contrast to TimeML, in the Richer Event Description³ (Styler et al., 2016, RED) framework the event/non-event distinction is not based on morpho-syntactic information. The RED Guidelines focus on those events (including occurrences, actions, processes and states) that deserve a place upon a timeline. Annotators should decide whether to mark something as an event or not on the basis of semantic questions regarding what is actually happening. According to these guidelines, *take (place)* in Example 1 above would not be annotated, since it merely helps constituting the description of the single ‘real’ event in this sentence expressed by *demonstrations*. In other words, only one event in this sentence should be put on a timeline. In this framework, syntax is only used to decide on the span of the event.

The two main corpora for factuality (or belief) are FactBank (Saurí and Pustejovsky, 2009) and the Lexical Understanding (LU) Annotation Corpus (Diab et al., 2009). Although both corpora address the same phenomenon (i.e. the commitment of a source towards the truth of some event/proposition), the annotations are quite different (Werner et al., 2015; Prabhakaran et al., 2015). The LU Corpus so far has only addressed the problem from the perspective of the speaker/writer, in contrast to FactBank, which has fully annotated nested sources. Furthermore, the LU corpus ignores negation (the polarity axis of factuality in FactBank) and does not distinguish between POSSIBLE and PROBABLE (the certainty axis of factuality in FactBank). Finally, there is a subtle difference in the targets. Whereas FactBank has assigned factuality values to events, the LU corpus has assigned belief tags to the head words of propositions, disregarding event-denoting noun phrases (e.g. *the collapse of the building*). An ongoing annotation

³RED is developed as a synthesis of the THYME-TimeML guidelines, the Stanford Event coreference guidelines and the Carnegie Mellon University Event coreference guidelines.

effort that is taking place at the Linguistic Data Consortium (LDC) aims to extend the annotations of the LU corpus (Prabhakaran et al., 2015). The plans include: (1) the annotation of nested beliefs in a similar way as was done in FactBank, (2) the extension of the definition of target propositions by using semantic representations (as opposed to using the propositional head) and including the heads of noun phrases, (3) the identification of entities as targets of beliefs (referred to as the notion of belief “aboutness”), and (4) combining belief with sentiment annotations.

The annotation scheme behind the Multi-Perspective Question Answering (MPQA) Opinion Corpus (Wiebe et al., 2005) is defined around the notion of private states, which in terms of their functional components are described as (internal) states of experiencers holding attitudes, optionally towards targets. A distinction is made between DIRECT SUBJECTIVE FRAMES, which represent explicit mentions of private states as well as speech events expressing private states, and EXPRESSIVE SUBJECTIVE ELEMENT FRAMES, which represent expressions that indirectly express private states through the way something is described or through a particular wording. Other frames that have been defined in MPQA are the OBJECTIVE SPEECH EVENT FRAME to distinguish opinion-oriented material from material presented as ‘factual’, the AGENT FRAME for representing the (nested) source of the attitude or speech event, the TARGET FRAME for representing the target of the opinion, and the ATTITUDE FRAME for characterizing the type (e.g. speculation, sentiment) and strength of the attitude.

Finally, the Penn Attribution Relations Corpus⁴ (Pareti, 2012, PARC) is a resource in which direct, indirect and mixed attributions of assertions, beliefs, facts and eventualities are annotated. In the PARC scheme, an attribution relation can consist of four elements: the cue, source, content, and supplement. The supplement is used to annotate additional information perceived as relevant for the interpretation of the attribution relation, such as recipients (e.g. *John told Mary that...*). In early versions of the PARC, attribution relations were further characterized by a set of features, including type (i.e. ASSERTION, BELIEF, FACT or EVENTUALITY), source (i.e. OTHER, ARBITRARY or WRITER), factuality (i.e. FACTUAL or NON-FACTUAL) and scopal change (i.e. SCOPAL CHANGE or NONE). However, due to low inter-annotator agreement, they are not present in the latest version (3.0) of the corpus.

When defining our annotation scheme, we aim to make our annotations as much compliant as possible with these major annotation initiatives on perspectives.

3. Fundamental Notions and Approach

Before we dive into the individual annotation layers, we introduce three fundamental parts of our approach. First, we describe the basic common elements of our annotations: source, cue and target. This is followed by a description of the GRaSP (Grounded Representation and Source Perspective) model, our formal representation of perspectives. Finally, we motivate our multilayered approach.

⁴The Penn Attribution Relations Corpus is an extension of the Penn Discourse TreeBank (Prasad et al., 2007).

3.1. Main Elements of Perspectives


Our perspective annotations are aimed at capturing the attitude of a source towards some target. The **source** can be any entity with an assumed individual or collective consciousness (e.g. a person, organization, or fictional character). The default source of a textual statement is its author,⁵ which is either implicitly present or lexicalized through first-person pronouns. Other possible sources are specific entities that are introduced in the text (e.g. *John believes that...*) or arbitrary sources expressed by non-specific references (e.g. *it is believed that...*).

An **attitude** can be described as a multidimensional characterization of the subjective relation between a source and target, which is a direct result of the beliefs and values of the source. At the moment of writing, we take two dimensions of attitude into account: factuality and opinion. Other dimensions (e.g. emotion) may be added in the future.

The **target** of a perspective can be an (abstract) entity, an event, or a (set of) propositional relation(s) expressing specific aspects of an event or entity. Consider the following examples, each of which illustrates a positive sentiment of *John* towards some target:

2. John likes Mary.
3. John is enjoying his birthday party organized by Mary.
4. John appreciates that Mary organized his birthday party.

In Example 2, the target of the positive sentiment is an entity referred to as *Mary*. In Example 3, the target is an event denoted by *his birthday party*. In Example 4, it is a propositional relation that could be represented as:



Note that both Examples 3 and 4 mention the ORGANIZEDBY relation between *Mary* and *his birthday party*; however, whereas in the former it functions merely as additional information on the target, in the latter it is the target of the perspective. The schematic representation in Figure 1 illustrates the differences between these targets. Not all attitude dimensions can take the same types of targets. For example, opinion can target entities, events or propositional relations, but factuality can only target events or propositional relations (Rambow and Wiebe, 2015).

Whereas the source and target are usually expressed by a single linguistic unit (e.g. an NP or a clause), the attitude may be expressed either by a single linguistic cue or a combination of cues.⁶ For example, the commitment of a source towards the factual nature of an event or proposition may be expressed by a combination of polarity (e.g. *not*, *never*) and modality cues (e.g. *could*, *maybe*). In turn, one cue can express multiple attitude dimensions. For example, a verb like *hope* expresses positive sentiment and uncertainty towards the target at the same time.

⁵If the author of a document is unknown, the default source is the document itself or its publisher.

⁶In fact, the selection of information itself is already considered perspective in our view.




Figure 1: Example of different perspective targets: entity, event and propositional relation

Despite the complexity of cues and the variation of targets and sources observed in different aspects of perspectives, we can generalize and define three main components of a perspective:

- **Source:** The lexical elements that refer to the entity which the perspective is attributed to.
- **Cue:** The lexical elements that signal the presence of a perspective and, individually or combined, characterize the relation between the source and the target (i.e. expressing the *attitude*).
- **Target:** The lexical elements that refer to the entity, event or (set of) propositional relation(s) functioning as the target of the perspective.

3.2. GRaSP

The annotations we describe in this paper can be integrated in a formal model called GRaSP (Grounded Representation and Source Perspective). GRaSP is an overarching model that provides the means to (1) represent **instances** (e.g. events, entities) and **propositions** in the (real or assumed) world, (2) to relate them to **mentions** in text⁷ using the Grounded Annotation Framework (Fokkens et al., 2013, GAF), and (3) to characterize the relation between mentions of sources and targets by means of **perspective-related annotations** such as attribution, factuality and sentiment.

We will illustrate GRaSP using the examples from Section 3.1. with an additional example, where the phrase *the event* refers to John’s birthday party as well:

5. Bill said Mary did not organize the event at all.

Figure 2 provides a simplified GRaSP representation of Sentences 4 and 5.⁸ The top part of Figure 2 represents instances mentioned in the sentences using the Simple Event Model (Van Hage et al., 2011, SEM). In this case, we have three SEM events: a party organized by Mary, an appreciation from John and a statement by Bill.

Bill and John have a different attitude to Mary organizing the party: John confirms and appreciates this and Bill denies it. These alternative perspectives displayed by the two

⁷In this paper we focus on textual mentions. However, in principal mentions can be anything referring to some instance: pictures, symbols, audio signals, etc.

⁸For reasons of clarity, the *say* and *appreciate* event representations are reduced and not all relations and attribution values are made explicit.

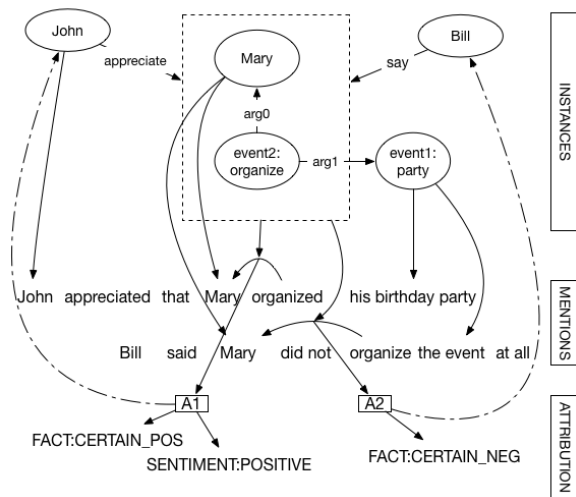


Figure 2: Simplified representation in GRASP

sentences are represented by GRASP as follows. There is one single representation of Mary organizing the party in the instance layer. The relation between Mary being the organizer is linked to their mentions in the two texts using GAF denotedBy relations (in black arrows from instance to mention layer). These relations make the fact that these mentions refer to Mary organizing the party explicit. Attitudes are represented in GRASP by linking source and target mentions by means of the GRASP relation *hasAttribution*. The relation between the source and target (i.e. the attitude) is characterized by the annotations in the different perspective layers, and is represented as a set of values identified by a unique attribution identifier. Figure 2 shows the attribution identifier of the relations (i.e. A1 and A2) between *Mary* and *organize(d)* in both sentences. These attribution identifiers can be linked to several perspective values: the factuality that each source assigns to it and the attributed sentiment, if present. The model also captures the source of the claim by linking attribution A1 to John and A2 to Bill (in dashed arrows).

The perspective annotations are done at the mention level using a multilayered annotation approach, which is described in the following subsection.

3.3. Multilayered Annotation

An important property of our annotation scheme is that it is multilayered. That is, we have defined separate layers for different semantic and pragmatic phenomena involved in the expression of perspectives, and these layers are annotated in a logical order. The need for a multilayered annotation scheme is motivated by two arguments.

First, our assumption is that human processing of texts naturally proceeds incrementally and compositionally by combining different pieces of information offered at the local level (e.g. by using lexical knowledge) to arrive at an interpretation at the more global level (e.g. by using context and world knowledge). Our multilayered annotation scheme is intended to simulate this process by starting with the local context (tokens, single clauses) and gradually moving to the global context (multiple clauses, multiple sentences).

Second, we believe that it is important to avoid presenting the annotators with an overload of information in the annotation process. Although at the unconscious level people

are perfectly capable of analyzing various dimensions of information in text, through previous experience we have found that it can become extremely confusing for annotators when they are asked to do the same at the conscious level. By splitting the annotations in subtasks, we avoid the problem of information overload caused by complex interactions between the different information layers involved. Instead, annotators can focus on one layer at a time, while still being able to use the annotations in the previous layers. In addition, it enables us to better monitor the interactions between the layers.

4. Layers of Perspective Annotation

In this section, we describe the four layers that we have currently defined for the annotation of perspectives: events, attribution, factuality, and opinion. All annotations are presented over the same sentence to show the interrelations of the different layers for representing perspectives.

4.1. Event Layer

In the first layer of our perspective annotation scheme we identify those lexical items that express **events**, which we consider to be the basic semantic units that may give rise to or be involved in perspectives. For this layer, we have considered both the TimeML and the RED guidelines. Since we are aiming at a more semantic approach to the annotation of perspectives, we have decided to use the RED guidelines as our basis for event annotations. According to the guidelines, a particular mention is only to be annotated if it constitutes its own event, not if it merely helps constituting the description of another event. This excludes, for example, grammaticalized verbs such as *occur* or *start*. This fits in well with our formal modeling of events and perspectives, where the information that is expressed by these kind of expressions is either not represented at all (because it does not add any semantic information), or represented as an (aspectual) relation with the main event. However, some adaptations to RED may be made if this is required for the integration with the other layers.

Following RED (and TimeML), we adhere to a minimal span annotation approach. However, following the News-Reader guidelines (Tonelli et al., 2014), we allow for some exceptions to the minimal span rule: the extent of phrasal verbs, idioms and prepositional phrases corresponds to the whole expression if they are entries either in the American or in the British version of the Collins English Dictionary online⁹ (e.g. *blow up*, *miss the boat*, *on a roll*).

In Example 6, all event mentions are marked in bold.

- Investors and Western diplomats have **said**_{e1} they might **interpret**_{e2} Mbeki's **support**_{e3} for Mugabe or the **elections**_{e4} as a sign that Africa is not intent on **revitalizing**_{e5} its economies through good **government**_{e6} and expanded international **trade**_{e7}.¹⁰

⁹<http://www.collinsdictionary.com/dictionary/english>

¹⁰MPQA Opinion Corpus – non_fbis-11.08.41-17418-S18

4.2. Attribution Layer

In the second layer we identify **attribution relations**. We follow the definition of attribution relations provided in PDTB and PARC, i.e. relations ascribing the ownership of an attitude towards some linguistic material. Our guidelines are based on and compliant with the PARC guidelines, though some adaptations are made to facilitate the integration with the other layers. The source, cue and target of this dimension can be specified as follows:

- **Source:** The entity that is the owner of the attributed abstract object, i.e. some content.
- **Cue:** The linguistic cue that signals the presence of an attribution relation and links the source to the target.
- **Target:** The abstract object that is being attributed to the source, being either a (textual) statement, a portion of it or its semantic content.

The source is either a reference to the author (e.g. *I believe that...*), a (vague or specific) third-party introduced in the text (e.g. *the minister/someone believes that...*), or an arbitrary source expressed by the use of non-specific references (e.g. *it is believed that...*).

The presence of an attribution relation can be signaled by different types of cues, including speech act verbs (e.g. *say, shout, think*), nouns (e.g. *statement, promise*), adjectives (e.g. *angry/sad that...*), prepositions or adverbs (e.g. *according to, reportedly*), or punctuation marks (Pareti, 2012). Some of them will already have been marked as events in the previous layer, in particular the verbs, nouns and adjectives referring to speech acts. In the annotation process, we do not create separate markables for these attribution cues, but rather annotate the event markable with an attribute indicating that the event is *attributional*.¹¹ Other cues will have to be annotated with a separate markable.

In our example, repeated in Example 7, there are several attribution relations, such as the one signaled by *said*. The spans of the source, cue and target of this relation are represented by means of curly brackets.

7. {Investors and Western diplomats}_{ATTR-SOURCE} have {**said**_{e1}}_{ATTR-CUE} {they might **interpret**_{e2} Mbeki's **support**_{e3} for Mugabe or the **elections**_{e4} as a sign that Africa is not intent on **revitalizing**_{e5} its economies through good **government**_{e6} and expanded international **trade**_{e7}}_{ATTR-TARGET}.

Inside the span of this target, there are other (nested) attribution relations signaled by *interpret* and *intent* (for clarity's sake, we did not represent nested attribution relations in the sentence). Although *support* does signal the ownership of an attitude (that of *Mbeki*), according to the PARC guidelines this should not be annotated as attribution since

¹¹According to the PARC guidelines, verbal attribution cues should be annotated together with their full verbal group, including auxiliaries, modals and negative particles. However, since we want to minimize the annotation effort and integrate the different layers, we adopt the minimal span rule of events for verbal cues.

the target (in PARC's terminology, the content) does not express linguistic material (a statement) or its semantic content. We do, however, annotate *support* as an attributional cue in the opinion layer (see Section 4.4.).

4.3. Factuality Layer

In the third layer we annotate **factuality** on top of the previous two layers. Our annotation approach is inspired by and compliant with FactBank (Saurí and Pustejovsky, 2009), where factuality is defined as the level of information expressing the commitment of relevant sources towards the factual nature of events mentioned in discourse. We define the three main elements of a factuality relation as follows:

- **Source:** The entity that commits to the factual nature of the targeted event.
- **Cue:** A linguistic cue that, possibly in combination with other cues, expresses the factual nature of the targeted event. We distinguish between the following cues:
 - *Attributional cue:* contributes a (new) source while expressing the commitment of this source towards the factual nature of the embedded event;
 - *Polarity cue (non-attributional):* affects the polarity of the event (affirmative or negative);
 - *Certainty cue (non-attributional):* indicates how certain the source is about the factual nature of the event.
- **Target:** The event which factual nature is being evaluated by the source.

The factuality layer is strongly connected to the event and attribution layers. Each event annotated in the event layer is taken as a target of a factuality relation. By default, the source of this relation is the author; if the event is embedded in an attribution relation, we create a factuality link between the event and the source markables annotated in the attribution layer through the attributional cue. In addition to the attributional cues already identified in the attribution layer, we annotate other relevant factuality cues as well, such as *might, a sign that* and *not* in our example sentence.

A source that is part of an attribution relation should be interpreted as a so-called nested source; that is, the position of the source towards the factuality of the event can only be learned through what the author asserts. In most cases, the authors themselves (or other sources that take a higher position in the nesting hierarchy) will remain uncommitted to the factuality of the event (Saurí and Pustejovsky, 2009). We deviate from FactBank's approach by only evaluating the factuality of relevant sources other than the *main source* (i.e. the source introduced by the attributional cue) if it can be understood that they clearly agree or disagree on the factual nature of the event.

A factuality relation is characterized by means of three attributes: certainty (CERTAIN, PROBABLE, POSSIBLE, UNDETERMINED), polarity (AFFIRMATIVE, NEGATIVE, UNDETERMINED) and time (FUTURE, NON-FUTURE, UNDETERMINED) (van Son et al., 2014). The combination of

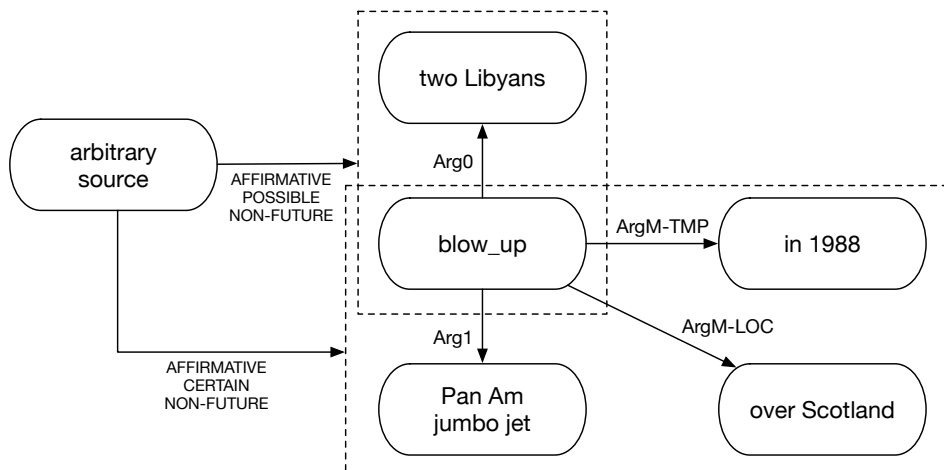


Figure 3: Schematic representation of an example of perspective scope (see Sentence 9)

the certainty and polarity attributes corresponds to the values used in FactBank, e.g. CT+ (CERTAIN/AFFIRMATIVE). The time attribute does not make the surface tense form value explicit, but is added to express the actual temporal interpretation of the annotated element. We believe this is a relevant attribute for factuality for two reasons. First, any future event does imply some degree of uncertainty, even when the source presents it with absolute certainty; therefore, we think it is relevant to make a distinction between past/present events and future events. Second, whether an event is placed in the past/present or in the future cannot always be derived from the tense form, as is the case for all non-verbal events, for example.

The factuality annotations of the target event *interpret* in our example sentence would be as follows:

8. {Investors and Western diplomats}_{ATTR-SOURCE} have {said_{e1}}_{ATTR-CUE} {they might interpret_{e2} Mbeki's support_{e3} for Mugabe or the elections_{e4} as a sign that Africa is not intent on revitalizing_{e5} its economies through good government_{e6} and expanded international trade_{e7}}_{ATTR-TARGET}.

Factuality annotations of “interpret” (e2):

Attributional cue: said
 Source (nested): {author, inv_dipl}
 Polarity cue: NA
 Certainty cue: might
 Factual assignments: POSSIBLE
 AFFIRMATIVE
 FUTURE

In the factuality layer, special attention is to be paid to event mentions for which we can derive any propositional relations from the sentence (for example, verbal events with one or more arguments, or nominal events including modifiers). For some of these mentions, it might not be sufficient to annotate the event, as a representative of the whole proposition, as the target of factuality. This is because factuality cues can target **specific relations** within a proposition. To clarify, consider the following example, taken from FactBank:

9. The World Court Friday **rejected**_{e1} U.S. and British **objections**_{e2} to a Libyan World Court **case**_{e3} that has **blocked**_{e4} the **trial**_{e5} of two Libyans **suspected**_{e6} of **blowing up**_{e7} a Pan Am jumbo jet over Scotland in 1988.¹²

In FactBank, the factuality value assigned to *blowing up* in this sentence is PR+ (“did probably happen”) for both relevant sources, namely the AUTHOR and the GEN_AUTHOR.¹³ However, the uncertainty expressed by *suspected* is not so much directed towards the proposition as a whole as it is towards a specific aspect of the event: the *two Libyans* being the ones who did it. In other words, both the AUTHOR and the GEN_AUTHOR do commit to the factual status of the Pan Am jumbo jet being blown up over Scotland in 1988, but they are not fully certain about *who* was responsible (see Figure 3). Similarly, in our main example sentence, *the investors and Western diplomats* do not particularly call into question whether Africa is intent on revitalizing its economies at all, but rather question the manner in which this will be done: *through good government and expanded international trade* (ArgM-MNR of *revitalize*), or not. We therefore argue that in cases where the attitude of a source is directed towards one of the arguments of an event it should appropriately be represented in the annotations. For example, the factuality of *blowing up* in Sentence 9 could be annotated as follows:

Factuality annotations of “blowing up” (e7):

Attributional cue: suspected
 Source (nested): {author, gen_author}
 Polarity cue: NA
 Certainty cue: NA
 Factual assignments:
 f(e7,Arg0) POSSIBLE AFF. NON-FUT.
 f(e7,Arg1) CERTAIN AFF. NON-FUT.
 f(e7,ArgM-TMP) CERTAIN AFF. NON-FUT.
 f(e7,ArgM-LOC) CERTAIN AFF. NON-FUT.

¹²TimeBank/FactBank – APW19980227.476-S1

¹³GEN_AUTHOR denotes a non-explicit generic source.

We call this phenomenon **perspective scope**, referring to those specific propositional relations associated with an event (or entity) that are affected by a perspective cue. It is strongly related to the scope and focus of negation as investigated by Blanco and Moldovan (2011), but to our knowledge its annotation has not been investigated before in the context of factuality (or sentiment). We believe that perspective scope is an important and innovative aspect of our annotation scheme, and our formal model GRaSP allows for the representation of separate factuality assignments for the event and its relations. In the near future, we will work out the details with respect to its annotation.

4.4. Opinion Layer

The final annotation layer that we have included in our scheme is that of **opinion** or **sentiment**. As our annotations are largely based on Wiebe et al. (2005) and Toprak et al. (2010), the three main elements are defined as follows:

- **Source:** The entity that has a positive or negative attitude towards some target.
- **Cue:** A linguistic cue that, possibly in combination with other cues, expresses the positive or negative attitude of the source towards the target. We regard cues as belonging to one of the following categories:
 - *Attributional cue:* contributes a source while directly expressing the positive or negative attitude of the source towards the embedded target;
 - *Indirect cue:* signals the positive or negative attitude of the source by the choice of words;
 - *Factual opinion cue:* refers to facts that can objectively verified, but that - in the given context - imply an evaluation of value.
- **Target:** The entity, event or (set of) propositional relation(s) that is positively or negatively evaluated by the source.

The annotation of opinion requires a different approach than that of factuality. In the factuality layer, each event identified in the event layer is to be annotated as the target of a factuality relation. In the opinion layer, annotators have no clear pre-defined targets; instead, they need to look for cues and understand the text in more detail.

The first thing to look for are attributional cues, since they are fairly easy to recognize and some of them will already have been identified in the attribution layer. An example of such an attributional cue is *support* in our example sentence repeated below, which expresses a positive attitude of *Mbeki*. Following Deng and Wiebe (2015), we aim to identify the specific entities and events that are the target of the opinion. In this case, there are two targets: the entity denoted by *Mugabe* and the event expressed by *elections*.

10. Investors and Western diplomats have **said**_{e1} they might **interpret**_{e2} {*Mbeki's*}_{SENT-SOURCE} {**support**_{e3}}_{SENT-CUE} for {*Mugabe*}_{SENT-TARGET} of the {**elections**_{e4}}_{SENT-TARGET} as a sign that Africa is not intent on **revitalizing**_{e5} its economies through good **government**_{e6} and expanded international **trade**_{e7}.

The other two types of cues are also present in our example. An example of a factual opinion cue is *expanded international* (targeting *trade*), and an example of an indirect cue is *good* (targeting *government*). However, maybe the strongest attitudes that can be understood from this sentence are, firstly, the negative attitude of the investors and Western diplomats towards Mugabe and, secondly, their negative attitude towards Africa's economic policy. Both attitudes require inference from the complex expression *Africa is not intent on revitalizing its economies through good government and expanded international trade*. We are investigating ways to annotate these complex expressions of opinion in a more systematic way and with a different approach with respect to existing proposals by taking into account relations in the predicate-argument structure and their interaction with previous annotation layers.

5. Conclusion and Future Work

From a linguistic point of view, perspectives expressed in texts can be extremely complex. Previous research has addressed the different aspects of perspectives independently, resulting in several datasets annotated with different types of information. In this paper we have presented an annotation scheme that integrates the main linguistic phenomena involved in perspectives. The methodology that we propose for perspective annotation relies on a multilayered approach to keep the tasks small, clear and simple. In our current annotation scheme, which is still under development, we have defined four layers: events, attribution, factuality, and opinion. The annotations can be integrated in a formal model called GRaSP, which allows us to represent this information at the instance and mention levels. This way, we can represent alternative perspectives on the same entity, event or proposition next to each other.

Once the annotation scheme is completed, our goal is to experiment with several approaches to annotate a corpus, including semi-automatic and manual approaches. As for the semi-automatic approaches, we will experiment with training annotation tools on the existing corpora for each layer, which should be possible because we have adopted existing annotation schemes wherever possible. As for the manual approach, we aim at experimenting with crowdsourcing. The fact that we have defined a multilayered approach will facilitate the deployment of such tasks with a crowd of non-trained annotators. Inel et al. (2013) show how different layers of perspectives can be sequenced in a crowdsourcing-task-workflow where each of the layers is a separate crowdsourcing task. We plan to use the CrowdTruth framework (Inel et al., 2014), which harnesses the unique ability of the crowd to provide a wide range of points of views, perspectives and opinions, by introducing a novel approach to capturing the disagreement between crowd annotators (Aroyo and Welty, 2014) as an indicator for annotation quality, language ambiguity and semantic similarity of target annotations.

Especially with respect to perspective annotation, which is notoriously complex and rich, it will be interesting to compare crowd annotation with expert annotation. Not only do we want to test whether the crowd can achieve the same quality of annotation for such a complex task, but also

whether the crowd annotation preserves the compositionality of the perspective for complex constructions, possibly including double negations or nested perspectives. In other words, does the crowd read the text in the same way as experts do? The results from crowdsourcing experiments might shed light on aspects of the perspectives that are not covered by the expert annotation scheme.

Our annotation guidelines and pilot annotations are publicly available at <https://github.com/vua-perspectives>.

6. Acknowledgements

This work was supported by the Amsterdam Academic Alliance Data Science (AAA-DS) Program Award to the UvA and VU Universities, by the Dutch National Science Foundation from the Spinoza-2013 price and by the European Union’s 7th Framework Programme via the NewsReader project (ICT-316404).

7. Bibliographical References

- Aroyo, L. and Welty, C. (2014). The three sides of CrowdTruth. *Human Computation*, 1:31–34.
- Blanco, E. and Moldovan, D. (2011). Semantic representation of negation using focus detection. In *Proc. of the 49th ACLT: HLT-Vol 1*, pages 581–589.
- Deng, L. and Wiebe, J. (2015). MPQA 3.0: Entity/event-level sentiment corpus. In *Proc. of NAACL-HLT*.
- Diab, M., Levin, L., Mitamura, T., Rambow, O., Prabhakaran, V., and Guo, W. (2009). Committed belief annotation and tagging. In *Proc. of the LAW*, pages 68–73, Suntec, Singapore.
- Fokkens, A., van Erp, M., Vossen, P., Tonelli, S., van Hage, W., Serafini, L., Sprugnoli, R., and Hoeksema, J. (2013). GAF: A grounded annotations framework for events. In *Proc. of the 1st Workshop on Events*, Atlanta, USA.
- Inel, O., Aroyo, L., Welty, C., and Sips, R. (2013). Domain-independent quality measures for Crowd Truth Disagreement. *Proc. of DeRiVE-2013*, pages 2–13.
- Inel, O., Khamkham, K., Cristea, T., Dumitrache, A., Rutjes, A., van der Ploeg, J., Romaszko, L., Aroyo, L., and Sips, R.-J. (2014). CrowdTruth: machine-human computation framework for harnessing disagreement in gathering annotated data. In *The Semantic Web–ISWC 2014*, pages 486–504. Springer.
- Pareti, S. (2012). A database of attribution relations. In *Proc. of LREC*, pages 3213–3217.
- Prabhakaran, V., By, T., Hirschberg, J., Rambow, O., Shaikh, S., Strzalkowski, T., Tracey, J., Arrigo, M., Basu, R., Clark, M., Dalton, A., Diab, M., Guthrie, L., Prokofieva, A., Strassel, S., Werner, G., Wiebe, J., and Wilks, Y. (2015). A new dataset and evaluation for belief/factuality. In *Proc. of the 4th *SEM*, Denver, USA.
- Prasad, R., Dinesh, N., Lee, A., J. A., and Webber, B. (2007). Attribution and its annotation in the Penn Discourse TreeBank. *Traitement Automatique des Langues*, 47(2):43–64.
- Pustejovsky, J., Castano, J. M., Ingria, R., Sauri, R., Gaizauskas, R. J., Setzer, A., Katz, G., and Radev, D. R. (2003a). TimeML: Robust specification of event and temporal expressions in text. In *Proc. of the 5th IWCS*, Tilburg, The Netherlands.
- Pustejovsky, J., Hanks, P., Sauri, R., See, A., Gaizauskas, R., Setzer, A., Radev, D., Sundheim, B., Day, D., Ferro, L., and Lazo, M. (2003b). The TIMEBANK corpus. In *Corpus Linguistics*, volume 2003, pages 647–656.
- Pustejovsky, J., Lee, K., Bunt, H., and Romary, L. (2010). ISO-TimeML: An International Standard for Semantic Annotation. In *Proc. of LREC*, Malta.
- Rambow, O. and Wiebe, J. (2015). Sentiment and belief: How to think about, represent, and annotate private states. In *Proc. of ACL-IJCNLP*.
- Sauri, R. and Pustejovsky, J. (2009). FactBank: a corpus annotated with event factuality. *Language Resources and Evaluation*, 43(3):227–268.
- Styler, W., Crooks, K., O’Gorman, T., and Hamang, M. (2016). Richer Event Description (RED) Annotation Guidelines v.1.7 (unpublished manuscript). Technical report.
- Tonelli, S., Sprugnoli, R., Speranza, M., and Minard, A.-L. (2014). NewsReader Guidelines for Annotation at Document Level NWR-2014-2. Technical report.
- Toprak, C., Jakob, N., and Gurevych, I. (2010). Sentence and Expression Level Annotation of Opinions in User-Generated Discourse. In *Proc. of 48th ACL*, Uppsala, Sweden.
- Van Hage, W. R., Malaisé, V., Segers, R., Hollink, L., and Schreiber, G. (2011). Design and use of the Simple Event Model (SEM). *Journal of Web Semantics*, 9(2):128–136.
- van Son, C., van Erp, M., Fokkens, A., and Vossen, P. (2014). Hope and Fear: Interpreting Perspectives by Integrating Sentiment and Event Factuality. In *Proc. of LREC*, pages 3857–3864, Reykjavik, Iceland.
- Verhagen, M., Gaizauskas, R., Schilder, F., Hepple, M., Moszkowicz, J., and Pustejovsky, J. (2009). The TempEval challenge: identifying temporal relations in text. *Language Resources and Evaluation*, 43(2):161–179.
- Werner, G., Prabhakaran, V., Diab, M., and Rambow, O. (2015). Committed belief tagging on the FactBank and LU corpora: A comparative study. In *Proc. of NAACL Workshop on Extra-propositional aspects of meaning in computational linguistics (ExProM)*, Denver, USA.
- Wiebe, J., Wilson, T., and Cardie, C. (2005). Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-03):165–210.