GAF: A Grounded Annotation Framework for Events

Antske Fokkens, Marieke van Erp, Piek Vossen
The Network Institute
VU University Amsterdam
antske.fokkens@vu.nl
marieke.van.erp@vu.nl
piek.vossen@vu.nl

Sara Tonelli
FBK
Trento, Italy
satonelli@fbk.eu

Willem Robert van Hage
SynerScope B.V.
Eindhoven, The Netherlands
willem.van.hage@synerscope.com

Luciano Serafini, Rachele Sprugnoli
FBK
Trento, Italy
serafini@fbk.eu
sprugnoli@fbk.eu

Jesper Hoeksema
The Network Institute
VU University Amsterdam
j.e.hoeksema@vu.nl

Abstract

This paper introduces GAF, a grounded annotation framework to represent events in a formal context that can represent information from both textual and extra-textual sources. GAF makes a clear distinction between *mentions* of events in text and their formal representation as *instances* in a **semantic layer**. Instances are represented by RDF compliant URIs that are shared across different research disciplines. This allows us to complete textual information with external sources and facilitates reasoning. The semantic layer can integrate any linguistic information and is compatible with previous event representations in NLP. Through a use case on earthquakes in Southeast Asia, we demonstrate GAF flexibility and ability to reason over events with the aid of extra-linguistic resources.

1 Introduction

Events are not only described in textual documents, they are also represented in many other non-textual sources. These sources include videos, pictures, sensors or evidence from data registration such as mobile phone data, financial transactions and hospital registrations. Nevertheless, many approaches to textual event annotation consider events as text-internal-affairs, possibly across multiple documents but seldom across different modalities. It follows from the above that event representation is not exclusively a concern for the NLP community. It also plays a major role in several other branches of information science such as knowledge representation and the Semantic Web, which have created their own models for representing events.

We propose a grounded annotation framework (GAF) that allows us to interconnect different ways of describing and registering events, including non-linguistic sources. GAF representations can be used to reason over the cumulated and linked sources of knowledge and information to interpret the often incomplete and fragmented information that is provided by each source. We make a clear distinction between *mentions* of events in text or any other form of registration and their formal representation as *instances* in a **semantic layer**.

Mentions in text are annotated using the Terence Annotation Format (Moens et al., 2011, TAF) on top of which the semantic layer is realized using Semantic Web technologies and standards. In this semantic layer, instances are denoted with Uniform Resource Identifiers (URIs). Attributes and relations are expressed according to the Simple Event Model (Van Hage et al., 2011, SEM) and other established ontologies. Statements are grouped in named graphs based on provenance and (temporal) validity, enabling the representation of conflicting information. External knowledge can be related to instances from a wide variety of sources such as those found in the Linked Open Data Cloud (Bizer et al., 2009a).

Instances in the semantic layer can optionally be linked to one or more mentions in text or to other sources. Because linking instances is optional, our representation offers a straightforward way to include information that can be inferred from text, such as implied participants or whether an event is part of a series that is not explicitly mentioned. Due to the fact that each URI is unique, it is clear that mentions connected to the same URI have a coreferential relation. Other relations between instances (participants, subevents, temporal relations, etc.) are represented explicitly in the semantic layer.

The remainder of this paper is structured as follows. In Section 2, we present related work and explain the motivation behind our approach. Section 3 describes the in-text annotation approach. Our semantic annotation layer is presented in Section 4. Sections 5-7 present GAF through a use case on earthquakes in Indonesia. This is followed by our conclusions and future work in section 8.

2 Motivation and Background

Annotation of events and of relations between them has a long tradition in NLP. The MUC conferences (Grishman and Sundheim, 1996) in the 90s did not explicitly annotate events and coreference relations, but the templates used for evaluating the information extraction tasks indirectly can be seen as annotation of events represented in newswires. Such events are not ordered in time or further related to each other. In response, Setzer and Gaizauskas (2000) describe an annotation framework to create coherent temporal orderings of events represented in documents using closure rules. They suggest that reasoning with text independent models, such as a calendar, helps annotating textual representations.

More recently, generic corpora, such as Propbank (Palmer et al., 2005) and the Framenet corpus (Baker et al., 2003) have been built according to linguistic principles. The annotations aim at properly representing verb structures within a sentence context, focusing on verb arguments, semantic roles and other elements. In ACE 2004 (Linguistic Data Consortium, 2004b), event detection and linking is included as a pilot task for the first time, inspired by annotation schemes developed for named entities. They distinguish between event mentions and the trigger event, which is the mention that most clearly expresses its occurrence (Linguistic Data Consortium, 2004a). Typically, agreement on the trigger event is low across annotators (around 55% (Moens et al., 2011)). Timebank (Pustejovsky et al., 2006b) is a more recent corpus for representing events and time-expressions that includes temporal relations in addition to plain coreference relations.

All these approaches have in common that they consider the textual representation as a closed world within which events need to be represented. This means that mentions are linked to a trigger event or to each other but not to an independent semantic representation. More recently, researchers started to annotate events across multiple documents, such as the EventCorefBank (Bejan and Harabagiu, 2010). Cross-document coreference is more challenging for establishing the trigger event, but it is in essence not different from annotating textual event coreference within a single document. Descriptions of events across documents may complement each other providing a more complete picture, but still textual descriptions tend to be incomplete and sparse with respect to time, place and participants. At the same time, the comparison of events becomes more complex. We thus expect even lower agreement in assigning trigger events across documents. Nothman et al. (2012) define the trigger as the first new article that mentions an event, which is easier than to find the clearest description and still report inter-annotator agreement of .48 and .73, respectively.

Recent approaches to automatically resolve event coreference (cf. Chambers and Jurafsky (2011a), Bejan and Harabagiu (2010)) use some background data to establish coreference and other relations between events in text. Background information, including resources, and models learned from textual data do not represent mentions of events directly but are useful to fill gaps of knowledge in the textual descriptions. They do not alter the model for annotation as such.

We aim to take these recent efforts one step further and propose a grounded annotation framework (GAF). Our main goal is to integrate information from text analysis in a **formal context** shared with researchers across domains. Furthermore, GAF is **flexible** enough to contain contradictory information. This is both important to represent sources that (partially) contradict each other and to combine alternative annotations or output of different NLP tools. Because conflicting information may be present, **provenance** of information is provided in our framework, so that we may decide which source to trust more or use it as a feature to decide which interpretation to follow. Different models of event representation exist that can contribute valuable information. Therefore our model is compliant with prior approaches regardless of whether they are manual or automatic. Finally, GAF makes a clear distinction between *instances* and *instance mentions* avoiding the problem of determining a trigger event. Additionally, it facilitates the integration of information from extra-textual sources and information that can be inferred from texts, but is not explicitly mentioned. Sections 5 to 7 will explain how we can achieve this with GAF.

3 The TERENCE annotation format

The TERENCE Annotation Format (TAF) is defined within the TERENCE Project[1] with the goal to include event mentions, temporal expressions and participant mentions in a single annotation protocol (Moens et al., 2011). TAF is based on ISO-TimeML (Pustejovsky et al., 2010), but introduces several adaptations in order to fit the domain of children's stories for which it was originally developed. The format has been used to annotate around 30 children stories in Italian and 10 in English.

We selected TAF as the basis for our in-text annotation for three reasons. First, it incorporates the (in our opinion crucial) distinction between *instances* and *instance mentions*. Second, it adapts some consolidated paradigms for linguistic annotation such as TimeML for events and temporal expressions and ACE for participants and participant mentions (Linguistic Data Consortium, 2005). It is thus compatible with other annotation schemes. Third, it integrates the annotation of event mentions, participants and temporal expressions into a unified framework. We will elaborate briefly on these properties below.

As mentioned, TAF makes a clear distinction between *instances* and *instance mentions*. Originally, this distinction only applied to nominal and named entities, similar to ACE (Linguistic Data Consortium, 2005), because children's stories can generally be treated as a closed world, usually presenting a simple sequence of events that do not corefer. Event coreference and linking to other sources was thus not relevant for this domain. In GAF, we extend the distinction between instances and instance mentions to events to model event coreference, link them to other sources and create a consistent model for all instances.

Children's stories usually include a small set of characters, event sequences (mostly in chronological order), and a few generic temporal expressions. In the TERENCE project, modeling characters in the stories is necessary. This requires an extension of TimeML to deal with event participants. Pustejovsky et al. (2006a) address the need to include arguments in TimeML annotations, but that proposal did not include specific examples and details on how to perform annotation (e.g., on the participants' attributes). Such guidelines were created for TAF.

The TAF annotation of *event mentions* largely follows TimeML in annotating tense, aspect, class, mood, modality and polarity and temporal expressions. However, there are several differences between TAF and TimeML. First, temporal expressions are not normalized into the ISO-8601 form, because most children's stories are not fixed to a specific date. In GAF, the normalization of expressions takes place in the semantic layer as these go beyond the scope of the text. As a result, temporal vagueness of linguistic expressions in text do not need to be normalized in the textual representation to actual time points and remain underspecified.[2]

In TAF, events and participant mentions are linked through a *has_participant* relation, which is defined as a directional, one-to-one relation from the event to the participant mentions. Only mentions corresponding to mandatory arguments of the events in the story are annotated. Annotators look up each verb in a reference dictionary providing information on the predicate-argument structure of each verb. This makes annotation easier and generally not controversial. However, this kind of information can be provided only by annotators having a good knowledge of linguistics.

All annotations are performed with the Celct An-

[1] ICT FP7 Programme, ICT-2010-25410, `http://www.terenceproject.eu/`

[2] Note that we can still use existing tools for normalization at the linguistic level: early normalizations can be integrated in the semantic layer alongside normalizations carried out at a later point.

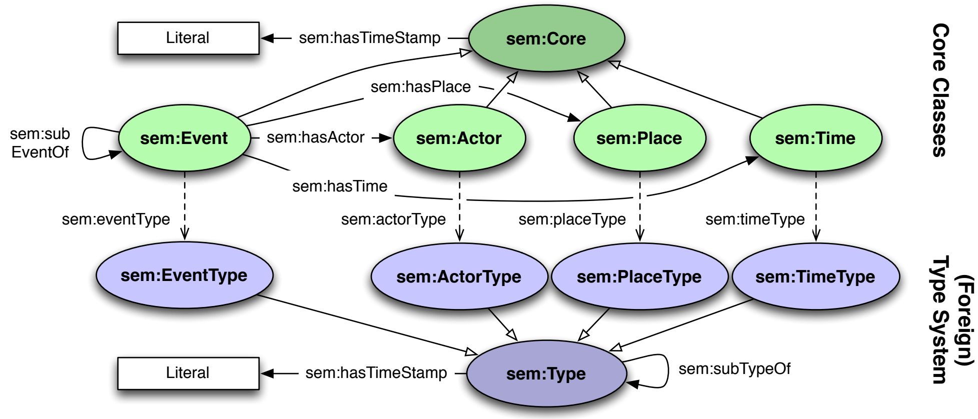

Figure 1: The SEM ontology

notation Tool (Bartalesi Lenzi et al., 2012), an online tool supporting TimeML that can easily be extended to include participant information. The annotated file can be exported to various XML formats and imported into the semantic layer. The next section describes SEM, the event model used in our semantic layer, and how it complements the TAF annotations.

4 The Simple Event Model

The Simple Event Model (SEM) is an RDF schema (Carroll and Klyne, 2004; Guha and Brickley, 2004) to express *who* did *what*, *where*, and *when*. There are many RDF schemas and OWL ontologies (Motik et al., 2009) that describe events, e.g., Shaw et al. (2009), Crofts et al. (2008) and Scherp et al. (2009). SEM is among the most flexible and easiest to adapt to different domains. SEM describes events and related instances such as the place, time and participants (called Actors in SEM) by representing the interactions between the instances with RDF triples. SEM models are semantic networks that include events, places, times, participants and all related concepts, such as their types.

An overview of all the classes in the SEM ontology and the relations connecting them is shown in Figure 1. Nodes can be identified by URIs, which universally identify them across all RDF models. If for example one uses the URI used by DBpedia[3] (Bizer et al., 2009b) for the 2004 catastrophe in Indonesia, then one really means the same event as everybody else who uses that URI. SEM does not put any constraints on the RDF vocabulary, so vocabularies can easily be reused. Places and place types can for example be imported from GeoNames[4] and event types from the RDF version of WordNet.

SEM supports two types of abstraction: generalization with hierarchical relations from other ontologies, such as the subclass relation from RDFS, and aggregation of events into superevents with the `sem:subEventOf` relation, as exemplified in Figure 2. Other types of abstractions can be represented using additional schemas or ontologies in combination with SEM. For instance, temporal aggregation can be done with constructs from the OWL Time ontology (Hobbs and Pan, 2004).

Relations between events and other instances, which could be other events, places, actors, times, or external concepts, can be modeled using the `sem:eventProperty` relation. This relation can be refined to represent specific relations, such as specific participation, causality or simultaneity relations. The provenance of information in the SEM graph is captured through assigning contexts to statements using the PROV Data Model (Moreau et al., 2012). In this manner, all statements derived from a specific newspaper article are stored in a named graph that represents that origin. Conflicting statements can be stored in different named graphs, and can thus coexist. This gives us the possibility

[3] `http://dbpedia.org`

[4] `http://www.geonames.org/ontology/`

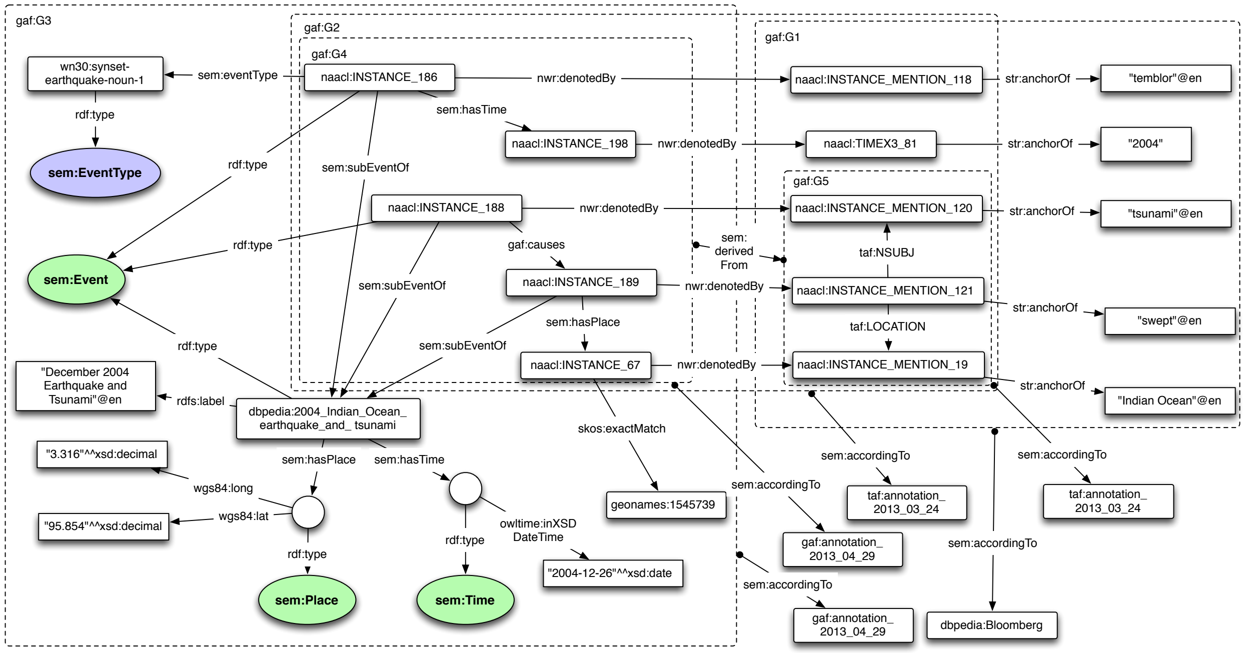

Figure 2: Partial SEM representation of December 26th 2004 Earthquake

of delaying or ignoring the resolution of the conflict, which enables use cases that require the analysis of the conflict itself.

5 The GAF Annotation Framework

This section explains the basic idea behind GAF by using texts on earthquakes in Indonesia. GAF provides a general model for event representation (including textual and extra-textual mentions) as well as exact representation of linguistic annotation or output of NLP tools. Simply put, GAF is the combination of textual analyses and formal semantic representations in RDF.

5.1 A SEM for earthquakes

We selected newspaper texts on the January 2009 West Papua earthquakes from Bejan and Harabagiu (2010) to illustrate GAF. This choice was made because the topic “earthquake” illustrates the advantage of sharing URIs across domains. Gao and Hunter (2011) propose a Linked Data model to capture major geological events such as earthquakes, volcano activity and tsunamis. They combine information from different seismological databases with the intention to provide more complete information to experts which may help to predict the occurrence of such events. The information can also be used in text interpretation. We can verify whether interpretations by NLP tools correspond to the data and relations defined by geologists or, through generalization, which interpretation is the most sensible given what we know about the events. General information on events obtained from automatic text processing, such as event templates (Chambers and Jurafsky, 2011b) or typical event durations (Gusev et al., 2010) can be integrated in SEM in a similar manner. Provenance indications can be used to indicate whether information is based on a model created by an expert or an automatically derived model obtained by a particular approach.

Figure 2 provides a fragment of a SEM representation for the earthquake and tsunami of December 26 2004.[5] The model is partially inspired by Gao and Hunter (2011)’s proposal. It combines information extracted from texts with information from DBpedia. The linking between the two can be established either manually or automatically through

[5]The annotation and a larger representation including the sentence it represents can be found on the GAF website `http://wordpress.let.vu.nl/gaf`.

an entity linking system.[6] The combined event of the earthquake and tsunami is represented by a DBpedia URI. The node labeled naacl:INSTANCE_186 represents the earthquake itself. The unambiguous representation of the 2004 earthquake leads us to additional information about it, for instance that the earthquake is an event (`sem:Event`) and that the `sem:EventType` is an earthquake, in this case represented by a synset from WordNet, but also the exact date it occurred and the exact location (cf `sem:hasTime, sem:hasPlace`).

5.2 Integrating TAF representations into SEM

TAF annotations are converted to SEM relations. For example, the TAF `as_participant` relations are translated to `sem:hasActor` relations, and temporal relations are translated to `sem:hasTime`. We use the relation `nwr:denotedBy` to link instances to their mentions in the text which are represented by their unique identifiers in Figure 2.

Named graphs are used to model the source of information as discussed in Section 4. The relation `sem:accordingTo` indicates provenance of information in the graph.[7] For instance, the mentions from the text in named graph `gaf:G1` come from the source `dbpedia:Bloomberg`. Relations between instances (e.g. between INSTANCE_189 and INSTANCE_188) are derived from a specific grammatical relation in the text (here, that *tsunami* is subject of *swept*) indicated by the `nwr:derivedFrom` relation from `gaf:G5` to `gaf:G4`. The grammatical relations included in graph `gaf:G5` come from a TAF annotation (`tag:annotation_2013_03_24`).

6 GAF Earthquake Examples

This section takes a closer look at a few selected sentences from the text that illustrate different aspects of GAF. Figure 2 showed how a URI can provide a **formal context** including important background information on the event. Several texts in the corpus refer to the tsunami of December 26, 2004, *a 9.1 temblor in 2004 caused a tsunami* and *The catastrophe four years ago*, among others. Compared to time expressions such as *2004* and *four years ago*, time indications extracted from external sources like DBpedia are not only more precise, but also permit us to correctly establish the fact that these expressions refer to the same event and thus indicate the same time. The articles were published in January 2009: a direct normalization of time indications would have placed the catastrophe in 2005. The **flexibility** to combine these seemingly conflicting time indications and delay normalization can be used to correctly interpret that *four years ago* early January 2009 refers to an event taking place at the end of December 2004.

A fragment relating to one of the earthquakes of January 2009: *The quake struck off the coast [...] 75 kilometers (50 miles) west of [....] Manokwari* provides a similar example. The expressions *75 kilometers* and *50 miles* are clearly meant to express the same distance, but not identical. The location is most likely neither exactly 75 km nor 50 miles. SEM can represent an underspecified location that is included in the correct region. The exact location of the earthquake can be found in external resources. We can include both distances as expressions of the location and decide whether they denote the general location or include the normalized locations as alternatives to those from external resources.

Different sources may report different details. Details may only be known later, or sources may report from a different perspective. As **provenance** information can be incorporated into the semantic layer, we can represent different perspectives, and choose which one to use when reasoning over the information. For example, the following phrases indicate the magnitude of the earthquakes that struck Manokwari on January 4, 2009:

the 7.7 magnitude quake (source: Xinhuanet)
two quakes, measuring 7.6 and 7.4 (source: Bloomberg)
One 7.3-magnitude tremor (source: Jakartapost)

The first two magnitude indicators (7.7, 7.6) are likely to pertain to the same earthquake, just as the second two (7.4, 7.3) are. Trust indicators can be found through the provenance trace of each men-

[6] Entity linking is the task of associating a mention to an instance in a knowledge base. Several approaches and tools for entity linking w.r.t. DBpedia and other data sets in the Linked Open Data cloud are available and achieve good performances, such as DBpedia Spotlight (Mendes et al., 2011); see (Rizzo and Troncy, 2011) for a comparison of tools.

[7] The use of named graphs in this way to denote context is compatible with the method used by Bozzato et al. (2012).

tion. Trust indicators can include the date on which it was published, properties of the creation process, the author, or publisher (Ceolin et al., 2010). Furthermore, because the URIs are shared across domains, we can link the information from the text to information from seismological databases, which may contain the exact measurement for the quake.

Similarly, external information obtained through shared links can help us establish coreference. Consider the sentences in Figure 3. There are several ways to establish that the same event is meant in all three sentences by using shared URIs and reasoning. All sentences give us approximate time indications, location of the affected area and casualties. Reasoning over these sentences combined with external knowledge allows us to infer facts such as that *undersea [...] off [...] Aceh* will be in the Indian Ocean, or that the affected countries listed in the first sentence are *countries around the Indian Ocean*, which constitutes the *Indian Ocean Community*. The number of casualties in combination of the approximate time indication or approximate location suffices to identify the earthquake and tsunami in Indonesia on December 26, 2004. The DBpedia representation contains additional information such as the magnitude, exact location of the quake and a list of affected countries, which can be used for additional verification. This example illustrates how a **formal context** using URIs that are shared across disciplines of information science can help to determine exact referents from limited or imprecise information.

7 Creating GAF

GAF entails integrating linguistic information (e.g. TAF annotations) into RDF models (e.g. SEM). The information in the model includes **provenance** that points back to specific annotations. There are two approaches to annotate text according to GAF. The first approach is bottom-up. Mentions are marked in the text as well as relations between them (participants, time, causal relations, basically anything except coreference). Consequently, these annotations are converted to SEM representations as explained above. Coreference is established by linking mentions to the same instance in SEM. The second approach is top-down. Here, annotators mark relations between instances (events, their participants, time relations, etc.) directly into SEM and then link these to mentions in the text.

As mention in Section 2, inter-annotator agreement on event annotation is generally low showing that it is challenging. The task is somewhat simplified in GAF, since it removes the problem of identifying an event trigger in the text. The GAF equivalent of the event trigger in other linguistic annotation approaches is an instance in SEM. However, other challenges such as which mentions to select are in principle not addressed by GAF, though differences in inter-annotator agreement may be found depending on whether the bottom-up approach or the top-down approach is selected. The formal context of SEM may help frame annotations, especially for domains such as earthquakes, where expert knowledge was used to create basic event models. This may help annotators while defining the correct relations between events. On the other hand, the top-down approach may lead to additional challenges, because annotators are forced to link events to unambiguous instances leading to hesitations as to when new instances should be introduced.

Currently, we only use the bottom-up approach. The main reason is the lack of an appropriate annotation tool to directly annotate information in SEM. We plan to perform comparative studies between the two annotation approaches in future work.

8 Conclusion and Future Work

We presented GAF, an event annotation framework in which textual mentions of events are grounded in a semantic model that facilitates linking these events to mentions in external (possibly non-textual) resources and thereby reasoning. We illustrated how GAF combines TAF and SEM through a use case on earthquakes. We explained that we aim for a representation that can combine textual and extra-linguistic information, provides a clear distinction between instances and instance mentions, is flexible enough to include conflicting information and clearly marks the provenance of information.

GAF ticks all these boxes. All instances are represented by URIs in a semantic layer following standard RDF representations that are shared across research disciplines. They are thus represented completely independent of the source and clearly distin-

```
There have been hundreds of earthquakes in Indonesia since a 9.1 temblor in 2004 caused a
tsunami that swept across the Indian Ocean, devastating coastal communities and leaving more
than 220,000 people dead in Indonesia, Sri Lanka, India, Thailand and other countries.
(Bloomberg, 2009-01-07 01:55 EST)

The catastrophe four years ago devastated Indian Ocean community and killed more than 230,000
people, over 170,000 of them in Aceh at northern tip of Sumatra Island of Indonesia.
(Xinhuanet, 2009-01-05 13:25:46 GMT)

In December 2004, a massive undersea quake off the western Indonesian province of Aceh
triggered a giant tsunami that left at least 230,000 people dead and missing in a dozen
countries facing the Indian Ocean.  (Aljazeera, 2009-01-05 08:49 GMT)
```

Figure 3: Sample sentences mentioning the December 2004 Indonesian earthquake from sample texts

guished from mentions in text or mentions in other sources. The Terence Annotation Format (TAF) provides a unified framework to annotate events, participants and temporal expressions (and the corresponding relations) by leaning on past, consolidated annotation experiences such TimeML and ACE. We will harmonize TAF, the Kyoto Annotation Format (Bosma et al., 2009, KAF) and the NLP Interchange Format (Hellmann et al., 2012, NIF) with respect to the textual representation in the near future. The NAF format includes the lessons learned from these predecessors: layered standoff representations using URI as identifiers and where possible standardized data categories. The formal semantic model (SEM) provides the flexibility to include conflicting information as well as indications of the provenance of this information. This allows us to use inferencing and reasoning over the cumulated and aggregated information, possibly exploiting the provenance of the type of information source. This flexibility also makes our representation compatible with all approaches dealing with event representation and detections mentioned in Section 2. It can include automatically learned templates as well as specific relations between events and time expressed in text. Moreover, it may simultaneously contain output of different NLP tools.

The proposed semantic layer may be simple, its flexibility in importing external knowledge may increase complexity in usage as it can model events in every thinkable domain. To resolve this issue, it is important to scope the domain by importing the appropriate vocabularies, but no more. When keeping this in mind, reasoning with SEM is shown to be rich but still versatile (Van Hage et al., 2012).

While GAF provides us with the desired granularity and flexibility for the event annotation tasks we envision, a thorough evaluation still needs to be carried out. This includes an evaluation of the annotations created with GAF compared to other annotation formats, as well as testing it within a greater application. A comparative study of top-down and bottom-up annotation will also be carried out. As already mentioned in Section 7, there is no appropriate modeling tool for SEM yet. We are currently using the CAT tool to create TAF annotations and convert those to SEM, but will develop a tool to annotate the semantic layer directly for this comparative study.

The most interesting effect of the GAF annotations is that it provides us with relatively simple access to a vast wealth of extra-linguistic information, which we can utilize in a variety of NLP tasks; some of the reasoning options that are made available by the pairing up with Semantic Web technology may for example aid us in identifying coreference relations between events. Investigating the implications of this combination of NLP and Semantic Web technologies lies at the heart of our future work.

Acknowledgements

We thank Francesco Corcoglioniti for his helpful comments and suggestions. The research leading to this paper was supported by the European Union's 7th Framework Programme via the NewsReader Project (ICT-316404) and by the BiographyNed project, funded by the Netherlands eScience Center (http://esciencecenter.nl/). Partners in BiographyNed are Huygens/ING Institute of the Dutch Academy of Sciences and VU University Amsterdam.

References

Collin F. Baker, Charles J. Fillmore, and Beau Cronin. 2003. The structure of the FrameNet database. *International Journal of Lexicography*, 16(3):281–296.

Valentina Bartalesi Lenzi, Giovanni Moretti, and Rachele Sprugnoli. 2012. CAT: the CELCT Annotation Tool. In *Proceedings of LREC 2012*.

Cosmin Bejan and Sandra Harabagiu. 2010. Unsupervised event coreference resolution with rich linguistic features. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1412–1422.

Christian Bizer, Tom Heath, and Tim Berners-Lee. 2009a. Linked data - the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22.

Christian Bizer, Jens Lehmann, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. 2009b. DBpedia - A crystallization point for the Web of Data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(3):154 – 165.

Wauter Bosma, Piek Vossen, Aitor Soroa, German Rigau, Maurizio Tesconi, Andrea Marchetti, Monica Monachini, and Carlo Aliprandi. 2009. KAF: a generic semantic annotation format. In *Proceedings of the 5th International Conference on Generative Approaches to the Lexicon GL 2009*, Pisa, Italy.

Loris Bozzato, Francesco Corcoglioniti, Martin Homola, Mathew Joseph, and Luciano Serafini. 2012. Managing contextualized knowledge with the ckr (poster). In *Proceedings of the 9th Extended Semantic Web Conference (ESWC 2012)*, May 27-31.

Jeremy J. Carroll and Graham Klyne. 2004. Resource description framework (RDF): Concepts and abstract syntax. W3C recommendation, W3C, February. http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/.

Davide Ceolin, Paul Groth, and Willem Robert Van Hage. 2010. Calculating the trust of event descriptions using provenance. *Proceedings Of The SWPM*.

Nathanael Chambers and Dan Jurafsky. 2011a. Template-based information extraction without the templates. In *Proceedings of ACL-2011*.

Nathanael Chambers and Dan Jurafsky. 2011b. Template-based information extraction without the templates. In *Proceedings of ACL-2011*, Portland, OR.

Nick Crofts, Martin Doerr, Tony Gill, Stephen Stead, and Matthew Stiff. 2008. Definition of the CIDOC Conceptual Reference Model. Technical report, ICOM/CIDOC CRM Special Interest Group. version 4.2.5.

Lianli Gao and Jane Hunter. 2011. Publishing, linking and annotating events via interactive timelines: an earth sciences case study. In *DeRiVE 2011 (Detection, Representation, and Exploitation of Events in the Semantic Web) Workshop in conjunction with ISWC 2011*, Bonn, Germany.

Ralph Grishman and Beth Sundheim. 1996. Message understanding conference - 6: A brief history. In *Proceedings of the 16th conference on Computational linguistics (COLING'96)*, pages 466–471.

Ramanathan V. Guha and Dan Brickley. 2004. RDF vocabulary description language 1.0: RDF schema. W3C recommendation, W3C, February. http://www.w3.org/TR/2004/REC-rdf-schema-20040210/.

Andrey Gusev, Nathanael Chambers, Pranav Khaitan, Divye Khilnani, Steven Bethard, and Dan Jurafsky. 2010. Using query patterns to learn the duration of events. In *Proceedings of ISWC 2010*.

Sebastian Hellmann, Jens Lehmann, and Sören Auer. 2012. NIF: An ontology-based and linked-data-aware NLP Interchange Format. Working Draft.

Jerry R Hobbs and Feng Pan. 2004. An ontology of time for the semantic web. *ACM Transactions on Asian Language Information Processing (TALIP)*, 3(1):66–85.

Linguistic Data Consortium. 2004a. Annotation Guidelines for Event Detection and Characterization (EDC). `http://projects.ldc.upenn.edu/ace/docs/EnglishEDCV2.0.pdf`.

Linguistic Data Consortium. 2004b. The ACE 2004 Evaluation Plan. Technical report, LDC.

Linguistic Data Consortium. 2005. ACE (Automatic Content Extraction) English annotation guidelines for entities. Version 6.6, July.

Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. Dbpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th International Conference on Semantic Systems*, I-Semantics '11, pages 1–8.

Marie-Francine Moens, Oleksandr Kolomiyets, Emanuele Pianta, Sara Tonelli, and Steven Bethard. 2011. D3.1: State-of-the-art and design of novel annotation languages and technologies: Updated version. Technical report, TERENCE project – ICT FP7 Programme – ICT-2010-25410.

Luc Moreau, Paolo Missier, Khalid Belhajjame, Reza B'Far, James Cheney, Sam Coppens, Stephen Cresswell, Yolanda Gil, Paul Groth, Graham Klyne, Timothy Lebo, Jim McCusker, Simon Miles, James Myers, Satya Sahoo, and Curt Tilmes. 2012. PROV-DM: The PROV Data Model. Technical report.

Boris Motik, Bijan Parsia, and Peter F. Patel-Schneider. 2009. OWL 2 Web Ontology

Language structural specification and functional-style syntax. W3C recommendation, W3C, October. `http://www.w3.org/TR/2009/REC-owl2-syntax-20091027/`.

Joel Nothman, Matthew Honnibal, Ben Hachey, and James R. Curran. 2012. Event linking: Grounding event reference in a news archive. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–232, Jeju Island, Korea, July. Association for Computational Linguistics.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106, 2013/03/12.

James Pustejovsky, Jessica Littman, and Roser Saurì. 2006a. Argument Structure in TimeML. In *Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum*.

James Pustejovsky, Jessica Littman, Roser Saurí, and Marc Verhagen. 2006b. Timebank 1.2 documentation. Technical report, Brandeis University, April.

James Pustejovsky, Kiyong Lee, Harry Bunt, and Laurent Romary. 2010. ISO-TimeML: An international standard for semantic annotation. In *Proceedings o the Fifth International Workshop on Interoperable Semantic Annotation*.

Giuseppe Rizzo and Raphaël Troncy. 2011. NERD: A framework for evaluating named entity recognition tools in the Web of data. In *Workshop on Web Scale Knowledge Extraction, colocated with ISWC 2011*.

Ansgar Scherp, Thomas Franz, Carsten Saathoff, and Steffen Staab. 2009. F–a model of events based on the foundational ontology dolce+ dns ultralight. In *Proceedings of the fifth international conference on Knowledge capture*, pages 137–144. ACM.

Andrea Setzer and Robert J. Gaizauskas. 2000. Annotating events and temporal information in newswire texts. In *LREC*. European Language Resources Association.

Ryan Shaw, Raphaël Troncy, and Lynda Hardman. 2009. LODE: Linking Open Descriptions of Events. In *4th Annual Asian Semantic Web Conference (ASWC'09)*, Shanghai, China.

Willem Robert Van Hage, Véronique Malaisé, Roxane Segers, Laura Hollink, and Guus Schreiber. 2011. Design and use of the simple event model (SEM). *Journal of Web Semantics*.

Willem Robert Van Hage, Marieke Van Erp, and Véronique Malaisé. 2012. Linked open piracy: A story about e-science, linked data, and statistics. *Journal on Data Semantics*, 1(3):187–201.