

Assigning Time-Stamps to Event-Clauses

Elena Filatova

Information Sciences Institute
University of Southern
California
elena@isi.edu

Eduard Hovy

Information Sciences Institute
University of Southern
California
hovy@isi.edu

Abstract

We describe a procedure for arranging into a time-line the contents of news stories describing the development of some situation. We describe the parts of the system that deal with 1. breaking sentences into event-clauses and 2. resolving both explicit and implicit temporal references. Evaluations show a performance of 52%, compared to humans.

1 Introduction

Linguists who have analyzed news stories (Schokkenbroek,1999; Bell,1997; Ohtsuka and Brewer,1992, etc.) noticed that “narratives¹ are about more than one event and these events are temporally ordered. Though it seems most logical to recapitulate events in the order in which they happened, i.e. in chronological order, the events are often presented in a different sequence”. The same paper states that “it is important to reconstruct the underlying event order² for narrative analysis to assign meaning to the sequence in which the events are narrated at the level of discourse structure....If the underlying event structure cannot be reconstructed, it may well be impossible to understand the narrative at all, let alone assign meaning to its structure”.

Several psycholinguistic experiments show the influence of event-arrangement in news stories on the ease of comprehension by readers. Duszak (1991) had readers reconstruct a news story from the randomized sentences. According to his experiments readers have a default strategy by which—in the absence of cues to the contrary—they re-impose chronological order on events in the discourse.

¹ Schokkenbroek (1999) uses the term *narrative* for news stories that relate more than one event.

² i.e., chronological order.

The problem of reconstructing the chronological order of events becomes more complicated if we have to deal with separate news stories, written at different times and describing the development of some situation, as is the case for multidocument summarization.

By judicious definition, one can make this problem easy or hard. Selecting only specific items to assign time-points to, and then measuring correctness on them alone, may give high performance but leave much of the text unassigned. We address the problem of assigning a time-point to *every* clause in the text.

Our approach is to break the news stories into their constituent events and to assign time-stamps—either time-points or time-intervals—to these events. When assigning time-stamps we analyze both implicit time references (mainly through the tense system) and explicit ones (temporal adverbials) such as ‘on Monday’, ‘in 1998’, etc. The result of the work is a prototype program which takes as input set of news stories broken into separate sentences and produces as output a text that combines all the events from all the articles, organized in chronological order.

2 Data

As data we used a set of news stories about an earthquake in Afghanistan that occurred at the end of May in 1998. These news stories were taken from CNN, ABC, and APW websites for the DUC-2000 meeting. The stories were all written within one week. Some of the texts were written on the same day. In addition to a description of the May earthquake, these texts contain references to another earthquake that occurred in the same region in February 1998.

3 Identifying Events

To divide sentences into event-clauses we use CONTEX (Hermjakob, 1997), a parser that

produces a syntactic parse tree augmented with semantic labels. CONTEX uses machine learning techniques to induce a grammar from a given treebanks. A sample output of CONTEX is given in Appendix 1.

To divide a sentence into event-clauses the parse tree output by CONTEX is analyzed from left to right (root to leaf). The *::CAT* field for each node provides the necessary information about whether the node under consideration forms a part of its upper level event or whether it introduces a new event. *::CAT* features that indicate new events are: S-CLAUSE, S-SNT, S-SUB-CLAUSE, S-PART-CLAUSE, S-REL-CLAUSE. These features mark clauses which contain both subject (one or several NPs) and predicate (VP containing one or several verbs).

The above procedure classifies a clause containing more than one verb as a simple clause. Such clauses are treated as one event and only one time-point will be assigned to them. This is fine when the second verb is used in the same tense as the first, but may be wrong in some cases, as in *He lives in this house now and will stay here for one more year*. There are no such clauses in the analyzed data, so we ignore this complication for the present.

The parse tree also gives information about the tense of verbs, used later for time assignment.

In order to facilitate subsequent processing, we wish to rephrase relative clauses as full independent sentences. We therefore have to replace pronouns where it is possible by their antecedents. Very often the parser gives information about the referential antecedents (in the example below, *Russia*). Therefore we introduced the rule: if it is possible to identify the referent, put it into the event-clause:

1. *Russia* <...> said;

2. **which** <*Russia*> has loaned helicopters in previous disasters;

3. **it** <*Russia*> would consider sending aid.

But sometimes the antecedent is identified incorrectly.

Qulle charged that the United Nations and non-governmental organizations involved in the relief were poorly coordinated, which was costing lives.

Here the antecedent for *which* is identified as *the relief*, and gives *which* <*the relief*> was costing lives instead of *which* <*poor coordination*> was

costing lives. Fortunately, in most cases our rule works correctly.

Although the event-identifier works reasonably well, breaking text into event-clauses needs further investigation. Table 1 shows the performance of the system. Two kinds of mistakes are made by the event identifier: those caused by CONTEX (it does not identify clauses with omitted predicate, etc.) and those caused by the fact that our clause identifier does too shallow analysis of the parse tree.

4 Time-stamper

According to (Bell, 1997) “time is expressed at different levels—in the morphology and syntax of the verb phrase, in time adverbials whether lexical or phrasal, and in the discourse structure of the stories above the sentence”.

4.1 Representation of Time-points and -intervals

For the present work we use slightly modified time representations suggested in (Allen, 1991). Formats used for time representation:

- {YYYY:DDD:W}³ Used when it is possible to point out the particular day the event occurred. {YYYY1:DDD1:W1},{YYYY2:DDD2:W2}...
- Used when it is possible to point out several concrete days when the events occurred.
- {YYYY1:DDD1:W1}---{YYYY2:DDD2:W2}
- Used when it is possible to point out a range of days when the event occurred.
- <<<<{YYYY:DDD:W} Used when it is possible to say the event occurred {YYYY:DDD:W} or earlier.
- >>>>{YYYY:DDD:W} Used when it is possible to say the event occurred {YYYY:DDD:W} or later.

4.2 Time-points Used for the Time-stamper

We use two anchoring time points:

1. Time of the article

We require that the first sentence for each article contains time information. For example:

T1 (05/30/1998:Saturday 18:35:42.49)
PAKINSTAN MAY BE PREPARING FOR
ANOTHER TEST.

The date information is in bold. We denote by ***T_i*** the reference time-point for the article, where *i*

³ YYYY—year number, DDD—absolute number of the day within the year (1–366), W—umber of the day in a week (1- Monday, ... 7- Saturday). If it is impossible to point out the day of the week then W is assigned 0.

$Recall = (\# \text{ of event-clauses correctly identified by system}) / (\# \text{ of event-clauses identified manually})$
 $Precision = (\# \text{ of event-clauses correctly identified by system}) / (\# \text{ of event-clauses identified by system})$

Text number	# of clauses by human	# of clauses by system	# correct	recall	precision
<i>Text 1</i>	7	6	5	$5/7 = 71.42\%$	$5/6 = 83.33\%$
<i>Text 2</i>	27	31	15	$15/27 = 55.55\%$	$15/31 = 48.38\%$
<i>Text 3</i>	5	8	3	$3/5 = 60\%$	$3/8 = 37.5\%$
<i>Text 4</i>	28	28	18	$18/28 = 64.28\%$	$18/28 = 64.28\%$
<i>Text 5</i>	33	36	19	$19/33 = 57.57\%$	$19/36 = 52.77\%$
<i>Text 6</i>	58	63	36	$36/58 = 62.07\%$	$36/63 = 57.14\%$
total	158	172	96	$96/158 = 60.76\%$	$96/172 = 55.81\%$

Table 1. Recall and precision scores for event identifier.

means that it is the time point of article *i*. The symbol T_i is used as a comparative time-point if the time the article was written is unknown. The information in brackets gives the *exact date* the article was written, which is the main anchor point for the time-stamper. The information about hours, minutes and seconds is ignored for the present.

2. Last time point assigned in the same sentence

While analyzing different event-clauses within the same sentence we keep track of what time-point was most recently assigned within this sentence. If needed, we can refer to this time-point. In case the most recent time information assigned is not a date but an interval we record information about both time boundaries. When the program proceeds to the next sentence, the variable for the most recently assigned date becomes undefined. In most cases this assumption works correctly (example 5.2–5.3):

5.2.1 *In the village of Kol, hundreds of people swarmed a United Nations helicopter*

5.2.2 *that <a United Nations helicopter> touched down three days after Saturday's earthquake*

5.2.3 *<after Saturday's earthquake> struck a remote mountainous area rocked three months earlier by another massive quake*

5.2.4 *that <another massive quake> claimed some 2,300 victims.*

5.3.1 *On Monday and Tuesday, U.N. helicopters evacuated 50 of the most seriously injured to emergency medical centers.*

The last time interval assigned for sentence 5.2 is {1998:53:0}---{1998:71:0}, which gives an approximate range of days when the previous

earthquake happened. But the information in sentence 5.3 is about the recent earthquake and not about the previous one of 3 months earlier, which is why it would be a mistake to point Monday and Tuesday within that range.

Mani and Wilson (2000) point out “over half of the errors [made by his time-stamper] were due to propagation of spreading of an incorrect event time to neighboring events”. The rule of dropping the most recently assigned date as an anchor point when proceeding to the next sentence very often helps us to avoid this problem.

There are however cases where dropping the most recent time as an anchor when proceeding to the next sentence causes errors:

4.8.1 *But in February a devastating earthquake in the same region killed 2,300 people and left thousands of people homeless.*

4.9.1 *At the time international aid workers suffered through a logistical nightmare to reach the snow-bound region with assistance.*

It is clear that sentence 4.9 is the continuation of sentence 4.8 and refers to the same time point (February earthquake). In this case our rule assigns the wrong time to 4.9.1. Still we retain this rule because it is more frequently correct than incorrect.

4.3 Preprocessing

First, the text divided into event-clauses is run through a program that extracts all the date-stamps (made available by Kevin Knight, ISI). In most cases this program does not miss any date-stamps and extracts only the correct ones. The only cases in which it did not work properly for the texts were:

- sentence 1.h1
PAKISTAN MAY BE PREPARING FOR ANOTHER TEST.

Here the modal verb *MAY* was assumed to be the month, given that it started with a capital letter.

- sentence 6.24
Tuberculosis is already common in the area where people live in close quarters and have poor hygiene

here the noun *quarters*, which in this case is used in the sense *immediate contact or close range* (Merriam-Webster dictionary), was assumed to be used in the sense *the fourth part of a measure of time* (Merriam-Webster dictionary).

After extracting all the date-phrases we proceed to time assignment.

4.4 Rules of Time Assignment

When assigning a time to an event, we select the time to be either the most recently assigned date or, if the value of the most recently assigned date is undefined, to the date of the article. We use a set of rules to perform this selection. These rules can be divided into two main categories: those that work for sentences containing explicit date information, and those that work for sentences that do not.

4.4.1 Assigning Time-Stamps to the Clauses with Explicit Date Information

- Day of the Week

If the day-of-the-week used in the event-clause is the same as that of the article (or the most recently assigned date, if it is defined), and there no words before it could signal that the described event happened earlier or will happen later, then the time-point of the article (or the most recently assigned date, if it is defined) is assigned to this event. If before or after a day-of-the-week there is a word/words signaling that the event happened earlier or will happen later then the time-point is assigned in accordance with this signal-word and the most recently assigned date, if it is defined.

If the day-of-the-week used in the event-clause is not the same as that of the article (or the most recently assigned date, if it is defined), then if there are words pointing out that the event happened before the article was written or the tense used in the clause is past, then the time for the event-clause is assigned in accordance with this word (such words we call signal-words), or

the most recent day corresponding to the current day-of-the-week is chosen. If the signal-word points out that the event will happen after the article was written or the tense used in the clause is future, then the time for the event-clause is assigned in accordance with the signal word or the closest subsequent day corresponding to the current day-of-the-week.

5.3.1 On Monday and Tuesday, U.N. helicopters evacuated 50 of the most seriously injured to emergency medical centers.

The time for article 5 is (06/06/1998:Tuesday 15:17:00). So, the time assigned to this event-clause is: 5.3.1 {1998:151:1}, {1998:152:2}.

- Name of Month

The rules are the same as for a day-of-the-week, but in this case a time-range is assigned to the event-clause. The left boundary of the range is the first day of the month, the right boundary is the last day of the month, and though it is possible to figure out the days of weeks for these boundaries, this aspect is ignored for the present.

4.8.1 But in February a devastating earthquake in the same region killed 2,300 people and left thousands of people homeless.

The time for article 4 is (05/30/1998:Saturday 14:41:00). So, the time assigned to this event-clause is 4.8.1 {1998:32:0}---{1998:60:0}.

In the analyzed corpus there is a case where the presence of a name of month leads to a wrong time-stamping:

6.3.1 Estimates say

6.3.2 up to 5,000 people died from the May 30 quake,

6.3.3 more than twice as many fatalities as in the February disaster.

Because of *February*, a wrong time-interval is assigned to clause 6.3.3, namely {1998:32:0}---{1998:60:0}. As this event-clause is a description of the latest news as compared to some figures it should have the time-point of the article. Such cases present a good possibility for the use of machine learning techniques to disambiguate between the cases where we should take into account date-phrase information and where not.

- Weeks, Days, Months, Years

We might have date-stamps where the words *weeks, days, months, years* are used with modifiers. For example

5.2.1 *In the village of Kol, hundreds of people swarmed a United Nations helicopter*

5.2.2 *that <a United Nations helicopter> touched down three days after Saturday's earthquake*

5.2.3 *after Saturday's earthquake struck a remote mountainous area rocked **three months earlier** by another massive quake*

5.2.4 *that <another massive quake> claimed some 2,300 victims.*

In event-clause 5.2.3 the expression *three months earlier* is used. It is clear that to get the time for the event it is not enough to subtract 3 months from the time of the article because the above expression gives an approximate range within which this event could happen and not a particular date. For such cases we invented the following rule:

Time=multiplier*length⁴; (in this case 3*30);
Day=DDD-Time; (for *years* Year=YYYY-Time)
Left boundary of the range=

Day-round (10%(Day));

(for *years* = Year - round(10%(Year)))

Right boundary of the range =

Day + round (10%(Day));

(for *years* = Year + round (10%(Year)))

For event 5.2.3 the time range will be {1998:53:0}---{1998:71:0} (the exact date of the article is {1998:152:2}).

If the modifier used with *weeks, days, months* or *years* is *several*, then the multiplier used in (1) is equal to 2.

- When, Since, After, Before, etc.

If an event-clause does not contain any date-phrase but contains one of the words 'when', 'since', 'after', 'before', etc., it might mean that this clause refers to an event, the time of which can be used as a reference point for the event under analysis. In this case we ask the user to insert the time for this reference event manually.

This rule can cause problems in cases where 'after' or 'before' are used not as temporal connectors but as spatial ones, though in the analyzed texts we did not face this problem.

4.4.2 Assigning Time-Stamps to the Clauses without Explicit Date Information

⁴ For *days*, length is equal to 1, *weeks*–7, *months*–30.

- Present/Past Perfect

If the current event-clause refers to a time-point in Present/Past Perfect tense, then an open-ended time-interval is assigned to this event. The starting point is unknown; the end-point is either the most recently assigned date or the time-point of the article.

- Future Tense

If the current event-clause contains a verb in future tense (one of the verbs 'shall', 'will', 'should', 'would', 'might' is present in the clause) then the open-ended time-interval assigned to this event-clause has the starting point at either the most recently assigned date or the date of the article.

- Other Tenses

Other tenses that can be identified with the help of CONTEX are *Present* and *Past Indefinite*. In the analyzed data all the verbs in Present Indefinite are given the most recently assigned date (or the date of the article). The situation with Past Indefinite is much more complicated and requires further investigation of more data. News stories usually describe the events that already took place at some time in the past, which is why even if the day when the event happened is not over, past tense is very often used for the description (this is especially noticeable for US news of European, Asian, African and Australian events). This means that very often an event-clause containing a verb in Past Indefinite Tense can be assigned the most recently assigned date (or the date of the article). It might prove useful to use machine learned rules for such cases.

- No verb in the event-clause

If there is no verb in the event-clause then the most recently assigned date (or the date of the article) is assigned to the event-clause.

4.5 Sources of Errors for Time-stamper

We ran the time-stamper program on two types of data: list of event-clauses extracted by the event identifier and list of event-clauses created manually. Tables 2 and 3 show the results. In the former case we analyzed only the correctly identified clauses. One can see that even on manually created data the performance of the time-stamper is not 100%. Why?

Some errors are caused by assigning the time based on the date-phrase present in the event-clause, when this date-phrase is not an adverbial time modifier but an attribute. For example,

1. *Estimates say*
2. *up to 5,000 people died from the May 30 earthquake,*
3. *more than twice as many fatalities as in the February disaster.*

The third event describes the May 30 earthquake but the time interval given for this event is {1998:32:0}---{1998:60:0} (i.e., the event happened in February). It might be possible to use machine learned rules to correct such cases.

One more significant source of errors is the writing style:

1. *“When I left **early this morning**,*
2. *everything was fine.*
3. *After the earthquake, I came back,*
4. *and the house had collapsed.*
5. *I **looked for two days** and gave up.*
6. *Everybody gave up*

When the reader sees *early this morning* he or she tends to assign to this clause the time of the article, but later as seeing *looked for two days*, realizes that the time of the clause containing *early this morning* is two days earlier than the time of the article. It seems that the errors caused by the writing style can hardly be avoided.

If an event happened at some time-point but according to the information in the sentence we can assign only a time-interval to this event (for example, *February Earthquake*) then we say that the time-interval is assigned correctly if the necessary time-point is within this time-interval

5 Time-line for Several News Stories and its Applications

After stamping all the news stories from the analyzed set, we arrange the event-clauses from all the articles into a chronological order. After doing that we obtain a new set of event-clauses which can easily be divided into two subsets—the first one containing all the references to the February earthquake, the second one containing the list of event-clauses in chronological order, describing what happened in May.

Such a text where all the events are organized in a chronological order might be very helpful in multidocument summarization, where it is important to include into the final summary not

only the most important information but also the most recent one. The output of the presented system gives the information about the time-order of the events described in several documents.

6 Related work

Several linguistic and psycholinguistic studies deal with the problem of time-arrangement of different texts. The research presented in these studies highlights many problems but does not solve them.

As for computational applications of time theories, most work was done on temporal expressions that appear in scheduling dialogues (Busemann et al., 1997; Alexandresson et al., 1997). There are many constraints on temporal expressions in this domain. The most relevant prior work is (Mani and Wilson, 2000), who implemented their system on news stories, introduced rules spreading time-stamps obtained with the help of explicit temporal expressions throughout the whole article, and invented machine learning rules for disambiguating between specific and generic use of temporal expressions (for example, whether *Christmas* is used to denote the 25th of December or to denote some period of time around the 25th of December). They also mention a problem of disambiguating between temporal expression and proper name, as in ‘USA Today’.

7 Conclusion

Bell (1997) notices “more research is needed on the effects of time structure on news comprehension. The hypothesis that the non-canonical news format does adversely affect understanding is a reasonable one on the basis of comprehension research into other narrative genres, but the degree to which familiarity with news models may mitigate these problems is unclear”. This research can greatly improve the performance of time-stamper and might lead to a list of machine learning rules for time detection.

In this paper we made an attempt to not just analyze and decode temporal expressions but to apply this analysis throughout the whole text and assign time-stamps to such type of clauses, which later could be used as separate sentences in various natural language applications, for example in multidocument summarization.

Text number	Number of event-clauses identified correctly	Number of time point correctly assigned to correctly identified clauses	percentage of correct assignment
<i>text 1</i>	5	4	80.00
<i>text 2</i>	15	15	100
<i>text 3</i>	3	2	66.67
<i>text 4</i>	18	17	94.44
<i>text 5</i>	19	17	89.47
<i>text 6</i>	36	24	66.66
Total	96	79	82.29

Table 2. Time-stamper performance on automatically claused texts (only correctly identified clauses are analyzed).

text number	number of manually created event-clauses	number of time point correctly assigned to manually created clauses	percentage of correct assignment
<i>target 1</i>	7	6	85.71
<i>target 2</i>	27	20	74.07
<i>target 3</i>	5	4	80.00
<i>target 4</i>	28	26	92.85
<i>target 5</i>	33	30	90.91
<i>target 6</i>	58	37	63.79
Total	158	123	77.85

Table 3. Time-stamper performance on manually (correct) claused texts.

References:

J. Alexandresson, N. Reithinger, and E. Maier, 1997. Insights into the Dialogue Processing of VERBMOBIL. *Proceedings of the Fifth Conference on Applied Natural Language Processing*, 33–40

J. Allen, 1991. Time and Time Again: The Many Ways to Represent Time. *Int'l. Journal of Intelligent Systems* 4, 341–356.

J. Allen and G. Ferguson, 1994. Actions and Events in Interval Temporal Logic. The University of Rochester, Technical Report 521

A. Bell, 1997. The Discourse Structure of News Structure. *Approaches to Media Discourse*, ed. A. Bell, 64–104.

S. Busemann, T. Decleek, A.K. Diagne, L. Dini, J. Klein, and S. Schmeier, 1997. Natural Language Dialogue Service for Appointment Scheduling Agents. *Proceedings of the Fifth Conference on Applied Natural Language Processing*, 25–32.

I. Mani and G. Wilson, 2000. Robust Temporal Processing of News. *Proceedings of the ACL-2000 Conference*.

C. Schokkenbroek, 1999. News Stories: Structure, Time and Evaluation. *Time & Society*, vol. 8(1), 59–98.

Appendix 1

CONTEX output for the sentence “Russia, which has loaned helicopters in previous disasters, said it would consider sending aid” (sentence 5.11) . The surface nodes (lexemes) are underlined>.

::NODE 1 ::SURF "Russia , which has loaned helicopters in previous disasters , said it would consider sending aid ." ::CAT S-SNT ::SUBS (2 20 21 29) ::LEX "say" ::FORMS (((MODE F-DECL) (PERSON F-THIRD-P) (NUMBER F-SING) (CASE F-NOM) (TENSE F-PAST-TENSE)))

::NODE 2 ::SURF "Russia , which has loaned helicopters in previous disasters ," ::CAT S-NP ::SUBS (3 7) ::PARENTS (1) ::LEX "Russia" ::ROLES (SUBJ) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P))) ::INDEX 2

::NODE 3 ::SURF "Russia ," ::CAT S-NP ::SUBS (4 6) ::PARENTS (2)::LEX "Russia" ::ROLES (PRED) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P))) ::INDEX 1

::NODE 4 ::SURF "Russia" ::CAT S-NP ::SUBS (5) ::PARENTS (3)::LEX "Russia" ::ROLES (PRED) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P)))

::NODE 5 ::SURF "Russia" ::CAT S-PROPER-NAME ::PARENTS (4) ::LEX "Russia" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 6 ::SURF "," ::CAT D-COMMA ::PARENTS (3) ::LEX "," ::ROLES (DUMMY)

::NODE 7 ::SURF "which has loaned helicopters in previous disasters ," ::CAT S-REL-CLAUSE ::SUBS (8 10 11 13 19) ::PARENTS (2) ::LEX "loan" ::ROLES (MOD) ::FORMS (((MODE F-DECL) (TENSE F-PERF-TENSE) (PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 8 ::SURF "which" ::CAT S-INTERR-NP ::SUBS (9) ::PARENTS (7) ::LEX "which" ::ROLES (SUBJ) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P))) ::PRON-REF 1

::NODE 9 ::SURF "which" ::CAT S-INTERR-PRON ::PARENTS (8)::LEX "which" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 10 ::SURF "has loaned" ::CAT S-VERB ::PARENTS (7) ::LEX "loan" ::ROLES (PRED) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P) (TENSE F-PERF-TENSE)))

::NODE 11 ::SURF "helicopters" ::CAT S-NP ::SUBS (12) ::PARENTS (7)::LEX "helicopter" ::ROLES (OBJ) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-PLURAL)))

::NODE 12 ::SURF "helicopters" ::CAT S-NOUN ::PARENTS (11)::LEX "helicopter" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-PLURAL)))

::NODE 13 ::SURF "in previous disasters" ::CAT S-PP ::SUBS (14 15) ::PARENTS (7) ::LEX "disaster" ::ROLES (LOCATION) ::FORMS (((NUMBER F-PLURAL) (PERSON F-THIRD-P)))

::NODE 14 ::SURF "in" ::CAT S-PREP ::PARENTS (13)::LEX "in" ::ROLES (P)

::NODE 15 ::SURF "previous disasters" ::CAT S-NP ::SUBS (16 18) ::PARENTS (13) ::LEX "disaster" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-PLURAL)))

::NODE 16 ::SURF "previous" ::CAT S-ADJP ::SUBS (17) ::PARENTS (15) ::LEX "previous" ::ROLES (MOD)

::NODE 17 ::SURF "previous" ::CAT S-ADJ ::PARENTS (16)::LEX "previous" ::ROLES (PRED)

::NODE 18 ::SURF "disasters" ::CAT S-NOUN ::PARENTS (15) ::LEX "disaster" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-PLURAL)))

::NODE 19 ::SURF "," ::CAT D-COMMA ::PARENTS (7) ::LEX "," ::ROLES (DUMMY)

::NODE 20 ::SURF "said" ::CAT S-VERB ::PARENTS (1)::LEX "say" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-SING) (TENSE F-PAST-TENSE)))

::NODE 21 ::SURF "it would consider sending aid" ::CAT S-SNT ::SUBS (22 24 25) ::PARENTS (1)::LEX "consider" ::ROLES (COMPL) ::FORMS (((MODE F-DECL) (GENDER F-NEUT) (PERSON F-THIRD-P) (NUMBER F-SING) (CASE F-NOM) (TENSE F-PRES-TENSE) (MODALS F-WOULD)))

::NODE 22 ::SURF "it" ::CAT S-NP ::SUBS (23) ::PARENTS (21)::LEX "PRON" ::ROLES (SUBJ) ::FORMS (((NUMBER F-SING) (PERSON F-THIRD-P) (GENDER F-NEUT)))::PRON-REF 2

::NODE 23 ::SURF "it" ::CAT S-REG-PRON ::PARENTS (22)::LEX "PRON" ::ROLES (PRED) ::FORMS (((GENDER F-NEUT) (PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 24 ::SURF "would consider" ::CAT S-TR-VERB ::PARENTS (21)::LEX "consider" ::ROLES (PRED) ::FORMS (((GENDER F-NEUT) (PERSON F-THIRD-P) (NUMBER F-SING) (TENSE F-PRES-TENSE) (MODALS F-WOULD)))

::NODE 25 ::SURF "sending aid" ::CAT S-INF-CLAUSE ::SUBS (26 27) ::PARENTS (21) ::LEX "send" ::ROLES (COMPL) ::FORMS (((TENSE F-PRES-PART) (MODE F-DECL)))

::NODE 26 ::SURF "sending" ::CAT S-DITR-VERB ::PARENTS (25)::LEX "send" ::ROLES (PRED) ::FORMS (((TENSE F-PRES-PART)))

::NODE 27 ::SURF "aid" ::CAT S-NP ::SUBS (28) ::PARENTS (25)::LEX "aid" ::ROLES (OBJ) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 28 ::SURF "aid" ::CAT S-COUNT-NOUN ::PARENTS (27)::LEX "aid" ::ROLES (PRED) ::FORMS (((PERSON F-THIRD-P) (NUMBER F-SING)))

::NODE 29 ::SURF "." ::CAT D-PERIOD ::PARENTS