Technical Report: Syntactically Encoded Semantic Relationships Type of Entailment Pairs

Elena Akhmatova
Centre for Language Technology
Macquarie University
Sydney, Australia
elena@ics.mq.edu.au

Mark Dras
Centre for Language Technology
Macquarie University
Sydney, Australia
madras@ics.mq.edu.au

1 Introduction

Most state-of-the-art approaches to RTE seek a generic approach to the task and do not differentiate between text-hypothesis pairs. However, a possible alternative is to consider subclasses of entailment pairs and build models to handle these specialties. An instance of this idea is proposed in (Vanderwende et al., 2005), where the complete set of entailment pairs is divided in two: those whose categorization could be accurately predicted based solely on syntactic cues and those where it is not the case. Their subsequent work (Vanderwende et al., 2006) presents an RTE system based on this their work.

The broader context of our work is to investigate their style of fragmentation. In this framework, a generic system would have additional special components that take care of the special subclasses of entailment pairs. Such a component is involved when a pair of its subclass is recognized.

This technical report describes the type of entailment pairs seperate recognition of which will be considered later as a component of such a system.

2 Syntactically encoded semantic relationships

We looked through the RTE2 test set and partitioned the set into several groups of entailments. Though the entailment pairs are different, for every word in the hypothesis there is often a word in the text from which it is entailed. It is not always so and we focus on the entailment pairs where this is not the case.

The entailment relationship we are focusing on can be named *an Entailment due to Syntactically Encoded Semantic Relationships (ESESR)*, as a specific syntactic construction in the text encodes a semantic relationship between its elements that is explicitly shown in the hypothesis.

Being more precise, the text-hypothesis pairs of interest have the following characteristics:

1. The hypothesis is a simple sentence. That is a sentence that consists of a subject, a predicate, and an object, and that contains no subordinate clauses.

2. Both subject and object of the hypothesis (or in some cases their morphological variants) are found in the text.

3. The predicate of the hypothesis has no match with anything in the text that is connected to the matches of the subject and the object of the hypothesis.

4. The matches of the subject and the object in the hypothesis can be connected to each other in the text by any syntactic relationship except being dependants of the same verb or a verb derivative.

Thus, the predicate of the hypothesis is the semantic relationship between its subject and object that is not explicitly defined in the text but is implicitly presented in the syntactic relationship between the matches of the subject and object of the hypothesis in the text. As we have already mentioned in condition 4 of the definition, the matches of the subject and the object in the hypothesis can be connected to

each other in the text almost by any syntactic relationship. The relationships that most often encode some semantics in them, are: apposition [1]; head and attribute in the noun phrase; or noun plus its prepositional attachment.

Following are examples of the entailments of the described type:

(1) *Text*: From Les Combes, in the Italian Alps, yesterday, where the Pope is on vacation, the Vatican's Press Office Director, Joaquin Navarro Valls, responded with a written statement to the accusations made by the Israeli government against Benedict XVI.

Hypothesis: Les Combes is located in the Italian Alps.

The location *Les Combes* is in the relation of apposition to *the Italian Alps*. This syntactic relation implicitly encodes the semantic relation represented by the words *is located* between the noun groups.

(2) *Text*: Lt. Jim Bowell of the Butler Township Fire Department said the 4:45 a.m. accident set fire to about 100 yards of woods.

Hypothesis: Jim Bowell is engaged by the Butler Township Fire Department.

Lt. Jim Bowell is connected syntactically to *the Butler Township Fire Department* via a preposition. That implicitly encodes a relation between the person and organization, *to be engaged by*.

(3) *Text*: Japan's Kyodo news agency said the US could be ready to set up a liaison office—the lowest level of diplomatic representation—in Pyongyang if it abandons its nuclear program.

Hypothesis: Kyodo news agency is based in Japan.

The attributive relationship between *Kyodo news agency* and *Japan* suggests but does not state explicitly the relationship *is based* between them. The Kyodo news agency *is based in* Japan is entailed from the attributive relationships between the nouns.

(4) *Text*: "Relative size and the power of the purse are certainly key factors," says Samuel L. Husk, executive director of the Council of Great City Schools.

Hypothesis: Samuel L. Husk *works for* the Council of Great City Schools.

Husk is *an executive director of the Council of Great City Schools*; the prepositional attachment *of the Council of Great City Schools* to a noun phrase *an executive director* suggests that *a director works for the Council of Great City Schools*, thus *Husk works for the Council of Great City Schools*.

(5) *Text*: The British ambassador to Egypt, Derek Plumbly, told Reuters on Monday that authorities had compiled the list of 10 based on lists from tour companies and from families whose relatives have not been in contact since the bombings.

Hypothesis: Derek Plumbly *resides in* Egypt.

Plumbly is *an ambassador to Egypt*, an ambassador resides the country he or she is an ambassador to. Thus, *the ambassador to Egypt* implicitly suggests that *Derek Plumbly resides in Egypt*.

(6) *Text*: Mugabe arrived in Beijing on Saturday and toured the northeastern province of Jilin, visiting the headquarters of First Automotive Works Group, China's top vehicle maker.

Hypothesis: First Automotive Works Group *is based in* China.

First Automotive Works Group is *China's top vehicle maker*, *China* is an attribute to *a vehicle maker* in the noun group *a China's top vehicle maker*, the attributive relation suggests the relation *is based* between *China* and *a vehicle maker*. Hence, the hypothesis is *First Automotive Works Group is based in China*.

3 Automatic classification

In this section we assess how well defined our subclass is. We have built an automatic classifier that marks entailment pairs as true if they are of the ES-ESR type and false otherwise. We first discuss the features used by the classifier, then the results of the classification using ML.

[1] We use the definition of (Quirket al., 1985) here

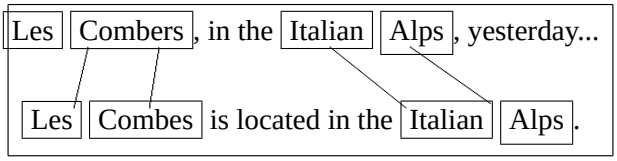

Figure 1: This is an example of word-to-word alignment of the text and hypothesis sentences

3.1 Features

As it is stated in the definition of the ESESR entailments the named entities in the text are most frequently related to each other through the relation of apposition, attribution or prepositional attachment. The same named entities are subject and direct or indirect object of the same predicate in the hypothesis. Thus, it makes sense to build two sets of syntactic features: ones that tell that the entailment is of a given type, and ones that most probably show that entailment is not of the given type.

To extract the features we build first the word-to-word alignment between the words of the text and hypothesis. Two words are aligned if they are similar or if they are not further from each other than on the path of the length 3 in WordNet (see figure 1). The decision on the length of the path was taken due to the time limitations for performing the parth search for all the pairs of words between the text sentence and the hypothesis sentence.

Features for syntactic correspondence between the aligned parts of text and hypothesis in favour of an entailment of given type are as below.

feature 1 The groups containing matches of the subject and the object of the hypothesis in text are in the relation of apposition (or preposition and noun phrase in commas follow the noun) to each other.

For example in (1) in Fig. 2, the connection of *the Italian Alps* and *Les Combes* in the text is similar to the relation of apposition. They are matches of *the Italian Alps* and *Les Combes* in the hypothesis respectively.

feature 2 The groups containing matches of the subject and the object of the hypothesis in the text are inside the same noun group.

For example in (2) in Fig. 2, *Japan* and *Kyodo news agency* are in the same noun group in the text. They match *Japan* and *the Kyodo news agency* in the hypothesis respectively.

feature 3 The groups containing matches of the subject and the object of the hypothesis in text represent a noun group and its prepositional attachment.

For example in (3) in Fig. 2, *Jim Bowell* is attached to *the Butler Township Fire Department* via a preposition *of*.

(1) *Text*: From Les Combes, in the Italian Alps, yesterday, where ....
Hypothesis: Les Combes is located in the Italian Alps.

(2) *Text*: Japan's Kyodo news agency ....
Hypothesis: The Kyodo news agency is based in Japan.

(3) *Text*: Lt. Jim Bowell of the Butler Township Fire Department ....
Hypothesis: Jim Bowell is engaged by the Butler Township Fire Department.

(4) *Text*: Samuel L. Husk, executive director of the Council of Great City Schools ....
Hypothesis: Samuel L. Husk works for the Council of Great City Schools.

(5) *Text*: Thompson's contract with the Dragons has been terminated after he reached an agreement with the club last night.
Hypothesis: The Dragons have terminated Thompson's contract.

Figure 2: Features for entailment classification

feature 4 The groups containing matches of the subject and the object of the hypothesis in text represent a chain of two syntactic connections, namely apposition relationship and a noun group plus its prepositional attachment.

For example in (4) in Fig. 2, *Samuel L. Husk* is connected by the relation of apposition to the *executive director* and *the executive director* in its turn continues the connection chain to *the Council of Great City Schools* by means of a preposition.

Syntactic correspondence between the aligned parts of text and hypothesis in favour of a pair not of this type (here we exploit the knowledge that matches of the subject and the object of the hypothesis are not dependants of verb or verb derivatives in text) are as below.

feature 5 Subject and predicate of the text are aligned to something in the hypothesis.

feature 6 Subject and object of the text are aligned to something in the hypothesis.

feature 7 Predicate and object of the text are aligned to something in the hypothesis.

As shown in the example (5) in Fig. 2, though *The Dragons* and *Thompson's contract* are aligned to *Thompson's contract with the Dragons*, *Thompson's contract has been terminated* is aligned to *have terminated Thompson's contract*. Thus whatever the hypothesis is about, it is not about the implicit relation in the *Thompson's contract with the Dragons*.

The selection of the semantic features is based on two facts. Firstly, the predicate of the hypothesis and only the predicate of the hypothesis is not aligned to anything in the text. Secondly, the subject and object of the hypothesis are usually named entities such as Person, Organization, Location, JobTitle.

feature 8 The percentage of the aligned words.

feature 9 The percentage of the not-aligned words.

In the example (5) of the figure 2 every word of the hypothesis corresponds to a word in the text, and the entailment pair is not of the type in question.

feature 10 A binary feature set to 1 if there are words in hypothesis that are not aligned with something in text.

feature 11 A binary feature set to 1 if the verb of hypothesis is aligned with something in text.

feature 12 The value of the feature is set to 1 whether there are named entities in both text and hypothesis.

feature 13 The value of the feature is set to 1 whether there are two or more named entities in both text and hypothesis.

feature 14 The value of the feature is equal to the number of matching named entities. Ideally this number is 2 for the ESESR entailment pairs, and 0 for non-ESESR entailment pairs.

There are also two features that combine sets of the syntactic features in one feature.

feature 15 The combination of the features 1, 2, 3 and 4. It is 0 if all of the features it combines are assigned a value of 0, and 1 if otherwise.

feature 16 The combination of the features 5, 6, 7. In a same manner as the previous feature, it is assigned to 0 if all the features it combines are assigned to 0, and 1 otherwise.

The last two features prove to be valuable as they are always among the top nodes in the decision trees, see figure 3.

3.2 Classification

The RTE2 test set consists of 800 entailment pairs. Only approximately one tenth of those pairs are ESESR entailments. To build the classifier we have duplicated all the ESESR entailment pairs several times to make the distribution of the entailment pairs equal. The reason for this is that we are interested in maximizing true positives, that it in a good recognition of the entailment pairs of the given class. Having only a small proportion of the set being of the ESESR type leads the machine learner to underweight these in the attempt to maximize the overall accuracy and gives a low TP rate, which is the one we are interested in. In this case we also had to make sure that taking the leaf out for the cross-validation assessment we take out all its copies as well. Otherwise the classifier is trained on the example we are

12 <= 0
— 9 <= 0: FALSE (342.0)
— 9 > 0
— — 13 <= 0
— — — 5 <= 0.474
— — — — 5 <= 0.2
— — — — — 14 <= 0: FALSE (4.0)
— — — — — 14 > 0: TRUE (17.0)
— — — — 5 > 0.2: FALSE (61.0)
— — — 5 > 0.474
— — — — 8 <= 0.182
— — — — — 6 <= 0: TRUE (7.0)
— — — — — 6 > 0
— — — — — — 5 <= 0.833: FALSE (44.0)
— — — — — — 5 > 0.833
— — — — — — — 10 <= 0
— — — — — — — — 16 <= 0
— — — — — — — — — 5 <= 0.857: TRUE (21.0/4.0)
— — — — — — — — — 5 > 0.857: FALSE (2.0)
— — — — — — — — 16 > 0: FALSE (5.0)
— — — — — — — 10 > 0: FALSE (8.0)
— — — — 8 > 0.182
— — — — — 14 <= 0
— — — — — — 10 <= 0
— — — — — — — 5 <= 0.727
— — — — — — — — 8 <= 0.3: FALSE (18.0)
— — — — — — — — 8 > 0.3
— — — — — — — — — 15 <= 0
— — — — — — — — — — 11 <= 0
— — — — — — — — — — — 5 <= 0.533: TRUE (18.0/1.0)
— — — — — — — — — — — 5 > 0.533: FALSE (6.0)
— — — — — — — — — — 11 > 0
— — — — — — — — — — — 7 <= 0
— — — — — — — — — — — — 8 <= 0.364: FALSE (5.0)
— — — — — — — — — — — — 8 > 0.364: TRUE (25.0/8.0)
— — — — — — — — — — — 7 > 0: TRUE (57.0/23.0)
— — — — — — — — — 15 > 0: FALSE (6.0)
— — — — — — — 5 > 0.727: TRUE (63.0/12.0)
— — — — — — 10 > 0
— — — — — — — 5 <= 0.727
— — — — — — — — 11 <= 1: TRUE (71.0/3.0)
— — — — — — — — 11 > 1
— — — — — — — — — 7 <= 0
— — — — — — — — — — 8 <= 0.3: TRUE (56.0/5.0)
— — — — — — — — — — 8 > 0.3: FALSE (4.0)
— — — — — — — — — 7 > 0
— — — — — — — — — — 8 <= 0.364: FALSE (7.0)
— — — — — — — — — — 8 > 0.364: TRUE (28.0/4.0)
— — — — — — — 5 > 0.727: FALSE (11.0)
— — — — — 14 > 0
— — — — — — 7 <= 0: FALSE (16.0)
— — — — — — 7 > 0
— — — — — — — 5 <= 0.5: TRUE (18.0/1.0)
— — — — — — — 5 > 0.5: FALSE (7.0)
— — 13 > 0
— — — 8 <= 0.583: FALSE (101.0)
— — — 8 > 0.583: TRUE (18.0/1.0)
12 > 0
— 10 <= 0
— — 4 <= 0
— — — 16 <= 0
— — — — 5 <= 0.583
— — — — — 7 <= 0: FALSE (9.0)
— — — — — 7 > 0
— — — — — — 9 <= 0: TRUE (19.0/2.0)
— — — — — — 9 > 0: FALSE (3.0)
— — — — 5 > 0.583: TRUE (97.0/12.0)
— — — 16 > 0: FALSE (8.0)
— — 4 > 0: FALSE (13.0)
— 10 > 0
— — 2 <= 0
— — — 1 <= 0
— — — — 11 <= 1: TRUE (104.0/2.0)
— — — — 11 > 1
— — — — — 5 <= 0.583: TRUE (41.0)
— — — — — 5 > 0.583: FALSE (5.0)
— — — 1 > 0: TRUE (105.0/1.0)
— — 2 > 0: TRUE (58.0)

Figure 3: J48 pruned tree (40 leaves, size 79) example

TP Rate	FP Rate	Precision	Recall	Class
0.87	0.39	0.69	0.87	FALSE
0.61	0.13	0.84	0.61	TRUE

Table 1: The result of the J48 classifier

going to classify that leads to the misleading results of the accuracy of the algorithm. We ran the J48 classifier on the dataset with the one-leaf-out cross validation test mode using the WEKA ML API (Witten and Frank, 1999). The overall accuracy is 75% (see table 1).

4 Conclusion

The outsome of the machine learner for the ESESR entailment type shows that the type is well-defined and can be isolated from all the other entailment pairs in order to apply a separate from the generic method for its recognition.

References

Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, Jan Svartvik 1985. A grammar of contemporary English, p 9-132 . *Longman, Singapore*.

Lucy Vanderwende, Deborah Coughlin, and William Dolan. 2006. What Syntax Can Contribute in the Entailment Task. In Proceedings of the PASCAL Challenges Workshop for Recognizing Textual Entailment. Southampton, U.K.

Lucy Vanderwende, Arul Menezes, and Rion Snow. 2006. Microsoft Research at RTE-2: Syntactic Contributions in the Entailment Task: an implementation. In Proceedings of the 2nd PASCAL Challenges Workshop on Recognising Textual Entailment. Venice, Italy

Ian H. Witten, Eibe Frank 1999. Practical machine learning tools with Java implementations. Morgan Kaufmann, San Francisco.