

A hybrid approach for extracting semantic relations from texts

Abstract

We present an approach for extracting relations from texts that exploits linguistic and empirical strategies, by means of a pipeline method involving a parser, part-of-speech tagger, named entity recognition system, pattern-based classification and word sense disambiguation models, and resources such as ontology, knowledge base and lexical databases. The relations extracted can be used for various tasks, including semantic web annotation and ontology learning. We suggest that the use of knowledge intensive strategies to process the input text and corpus-based techniques to deal with unpredicted cases and ambiguity problems allows to accurately discover the relevant relations between pairs of entities in that text.

1 Introduction

Semantic relations extracted from texts are useful for several applications, including question answering, information retrieval, semantic web annotation, and construction and extension of lexical resources and ontologies. In this paper we present an approach for relation extraction developed to semantically annotate relational knowledge coming from raw text, within a framework aiming to automatically acquire high quality semantic metadata for the Semantic Web.

In that framework, applications such as semantic web portals (Lei et al., 2006) analyze data from texts, databases, domain ontologies, and knowledge bases in order to extract the semantic knowledge in an integrated way. Known entities occurring in the text, i.e., entities that are included in the knowledge base, are semantically annotated with their properties, also provided by the knowledge base and by databases. New enti-

ties, as given by a named entity recognition system according to the possible types of entities in the ontology, are annotated without any additional information. In this context, the goal of the relation extraction approach presented here is to extract relational knowledge about entities, i.e., to identify the semantic relations between pairs of entities in the input texts. Entities can be both known and new, since named entity recognition is also carried out. Relations include those already existent in the knowledge base, new relations predicted as possible by the domain ontology, or completely new (unpredicted) relations.

The approach makes use of a domain ontology, a knowledge base, and lexical databases, along with knowledge-based and empirical resources and strategies for linguistic processing. These include a lemmatizer, syntactic parser, part-of-speech tagger, named entity recognition system, and pattern matching and word sense disambiguation models. The input data used in the experiments with our approach consists of English texts from the Knowledge Media Institute (KMi)¹ newsletters. We believe that by integrating corpus and knowledge-based techniques and using rich linguistic processing strategies in a completely automated fashion, the approach can achieve effective results, in terms of both accuracy and coverage.

With relational knowledge, a richer representation of the input data can be produced. Moreover, by identifying new entities, the relation extraction approach can also be applied to ontology population. Finally, since it extracts new relations, it can also be used as a first step for ontology learning.

In the remaining of this paper we first describe some cognate work on relation extraction, particularly those exploring empirical methods, for various applications (Section 2). We then present

¹ <http://kmi.open.ac.uk/>

our approach, showing its architecture and describing each of its main components (Section 3). Finally, we present the next steps (Section 4).

2 Related Work

Several approaches have been proposed for the extraction of relations from unstructured sources. Recently, they have focused on the use of supervised or unsupervised corpus-based techniques in order to automate the task. A very common approach is based on pattern matching, with patterns composed by subject-verb-object (SVO) tuples. Interesting work has been done on the unsupervised automatic detection of relations from a small number of seed patterns. These are used as a starting point to bootstrap the pattern learning process, by means of semantic similarity measures (Yangarber, 2000; Stevenson, 2004).

Most of the approaches for relation extraction rely on the mapping of syntactic dependencies, such as SVO, onto semantic relations, using either pattern matching or other strategies, such as probabilistic parsing for trees augmented with annotations for entities and relations (Miller et al, 2000), or clustering of semantically similar syntactic dependencies, according to their selectional restrictions (Gamallo et al., 2002).

In corpus-based approaches, many variations are found concerning the machine learning techniques used to produce classifiers to judge relation as relevant or non-relevant. (Roth and Yih, 2002), e.g., use probabilistic classifiers with constraints induced between relations and entities, such as selectional restrictions. Based on instances represented by a pair of entities and their position in a shallow parse tree, (Zelenko et al., 2003) use support vector machines and voted perceptron algorithms with a specialized kernel model. Also using kernel methods and support vector machines, (Zhao and Grishman, 2005) combine clues from different levels of syntactic information and applies composite kernels to integrate the individual kernels.

Similarly to our proposal, the framework presented by (Iria and Ciravegna, 2005) aims at the automation of semantic annotations according to ontologies. Several supervised algorithms can be used on the training data represented through a canonical graph-based data model. The framework includes a shallow linguistic processing step, in which corpora are analyzed and a representation is produced according to the data model, and a classification step, where classifiers run on the datasets produced by the linguistic

processing step.

Several relation extraction approaches have been proposed focusing on the task of ontology learning (Reinberger and Spyns, 2004; Schutz and Buitelaar, 2005; Ciaramita et al., 2005). More comprehensive reviews can be found in (Maedche, 2002) and (Gomez-Perez and Manzano-Macho, 2003). These approaches aim to learn non-taxonomic relations between concepts, instead of lexical items. However, in essence, they can employ similar techniques to extract the relations. Additional strategies can be applied to determine whether the relations can be lifted from lexical items to concepts, as well as to determine the most appropriate level of abstraction to describe a relation (e.g. Maedche, 2002).

In the next section we describe our relation extraction approach, which merges features that have shown to be effective in several of the previous works, in order to achieve more comprehensive and accurate results.

3 A hybrid approach for relation extraction

The proposed approach for relation extraction is illustrated in Figure 1. It employs knowledge-based and (supervised and unsupervised) corpus-based techniques. The core strategy consists of mapping linguistic components with some syntactic relationship (a linguistic triple) into their corresponding semantic components. This includes mapping not only the relations, but also the terms linked by those relations. The detection of the linguistic triples involves a series of linguistic processing steps. The mapping between terms and concepts is guided by a domain ontology and a named entity recognition system. The identification of the relations relies on the knowledge available in the domain ontology and in a lexical database, and on pattern-based classification and sense disambiguation models.

The main goal of this approach is to provide rich semantic annotations for the Semantic Web. Other potential applications include:

- 1) Ontology population: terms are mapped into new instances of concepts of an ontology, and relations between them are identified, according to the possible relations in that ontology.

- 3) Ontology learning: new relations between existent concepts are identified, and can be used as a first step to extend an existent ontology. A subsequent step to lift relations between instances to an adequate level of abstraction may be necessary.

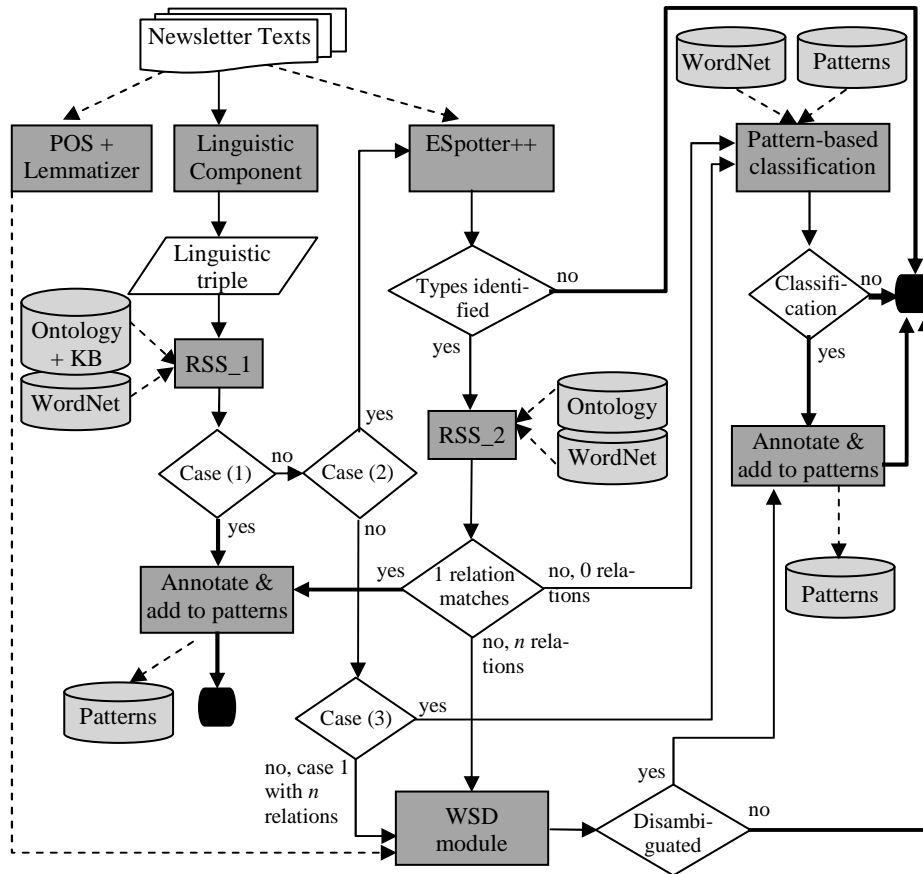


Figure 1. Architecture of the proposed approach

3.1 Context and resources

The input to our experiments consists of electronic **Newsletter Texts**². These are short texts describing news of several natures related to members of a research group: projects, publications, events, awards, etc. The domain **Ontology** used (*KMi-basic-portal-ontology*) was designed based on the AKT reference ontology³ to include concepts relevant to our domain. The instantiations of concepts in this ontology are stored in the knowledge base (**KB**) *KMi-basic-portal-kb*. The other two resources used in our architecture are the lexical database **WordNet** (Fellbaum, 1998) and a repository of **Patterns** of relations, described in Section 3.4.

3.2 Identifying linguistic triples

Given a newsletter text, the first step of the relation extraction approach is to process the natural language text in order to identify linguistic triples, that is, sets of three elements with a syntactic relationship, which can indicate potentially relevant semantic relations. In our architecture,

this is accomplished by the **Linguistic Component** module, and adaptation of the linguistic component designed in Aqualog (Lopez et al., 2005), a question answering system.

The linguistic component uses the infrastructure and the following resources from GATE (Cunningham et al., 2002): tokenizer, sentence splitter, part-of-speech tagger, morphological analyzer and VP chunker. On the top of these resources, which produce syntactic annotations for the input text, the linguistic component uses a grammar to identify linguistic triples. This grammar was implemented in Jape (Cunningham et al., 2000), which allows the definition of patterns to recognize regular expressions using the annotations provided by GATE.

The main type of construction aimed to be identified by our grammar involves a verbal expression as indicative of a potential relation and two noun phrases as terms linked by that relation. However, our patterns also account for other types of constructions, including, e.g., the use of comma to implicitly indicate a relation, as in sentence (1). In this case, when mapping the terms into entities (Section 3.3), having identified that “KMi” is an *organization* and “Enrico

² <http://news.kmi.open.ac.uk/kmiplanet/>

³ <http://kmi.open.ac.uk/projects/akt/ref-onto/>

Motta” is a *person*, it is possible to guess the relation indicated by the comma (e.g., work). Some examples triples identified by our patterns for the newsletter in Figure 2 are given in Figure 3.

(1) “Enrico Motta, at KMi now, is leading a project on”.

Nobel Summit on ICT and public services
 Peter Scott attended the Public Services Summit in Stockholm, during Nobel Week 2005. The theme this year was Responsive Citizen Centered Public Services. The event was hosted by the City of Stockholm and Cisco Systems Thursday 8 December - Sunday 11 December 2005.
 ...

Figure 2. Example of newsletter

<peter-scott,attend,public-services-summit>
 <public-services-summit,located,stockholm>
 <theme,is,responsive-citizen-centered-public-services>
 <city-of-stockholm-and-cisco-systems,host,event>

Figure 3. Examples of linguistic triples for the newsletter in Figure 2

Jape patterns are based on shallow syntactic information only, and therefore they are not able to capture certain potentially relevant triples. To overcome this limitation, we employ a parser as a complementary resource to produce linguistic triples. We use Minipar (Lin, 1993), which produces functional relations for the components in a sentence, including subject and object relations with respect to a verb. This allows capturing some implicit relations, such as indirect objects and long distance dependence relations.

Minipar’s representation is converted into a triple format and therefore the intermediate representation provided by both GATE and Minipar consists of triples of the type: <noun_phrase, verbal_expression, noun_phrase>.

3.3 Identifying entities and relations

Given a linguistic triple, the next step is to verify whether the verbal expression in that triple conveys a relevant semantic relationship between entities (given by the terms) potentially belonging to an ontology. This is the most important phase of our approach and is represented by a series of modules in our architecture in Figure 1. As first step we try to map the linguistic triple into an ontology triple, by using an adaptation of Aqualog’s Relation Similarity Service (RSS).

RSS tries to make sense of the linguistic triple by looking at the structure of the domain ontology and the information stored in the KB. In order to map a linguistic triple into an ontology triple, besides looking for an exact matching be-

tween the components of the two triples, RSS considers partial matching by using a set of resources in order to account for minor lexical or conceptual discrepancies between these two elements. These resources include metrics for string similarity matching, synonym relations given by WordNet, and a lexicon of previous mappings between the two types of triples. Different strategies are employed to identify a matching for terms and relations, as we describe below.

Since we do not consider any interaction with the user in order to achieve a fully automated annotation process, other modules were developed to complete the mapping process even if there is no matching (Section 3.4) or if there is ambiguity (Section 3.5), according to RSS.

Strategies for mapping terms

To map terms into entities, the following attempts are accomplished (in the given order):

1) Search the KB for an exact matching of the term with any instance.

2) Apply string similarity metrics⁴ to calculate the similarity between the given term and each instance of the KB. A hybrid scheme combining three metrics is used: jaro-Winkler, jlevelDistance a wlevelDistance. Different combinations of threshold values for the metrics are considered. The elements in the linguistic triples are lemmatized in order to avoid problems which could be incorrectly handled by the string similarity metrics (e.g., past tense).

2.1) If there is more than one possible matching, check whether any of them is a substring of the term. For example, the instance name for “Enrico Motta” is a substring of the term “Motta”, and thus it should be preferred.

For example, the similarity values returned for the term “vanessa” with instances potentially relevant for the mapping are given in Figure 4. The combination of thresholds is met for the instance “Vanessa Lopez”, and thus the mapping is accomplished. If there is still more than one possible mapping, we assume there is not enough evidence to map that term and discard the triple.

jaroDistance for “vanessa” and “vanessa-lopez” = 0.8461538461538461
 wlevel for “vanessa” and “vanessa-lopez” = 1.0
 jWinklerDistance for “vanessa” and “vanessa-lopez” = 0.9076923076923077

Figure 4. String similarity measures for the term “vanessa” and the instance “Vanessa Lopez”

⁴ <http://sourceforge.net/projects/simmetrics/>

Strategies for mapping relations

In order to map the verbal expression into a conceptual relation, we assume that the terms of the triple have already been mapped either into instances of classes in the KB by RSS, or into potential new instances, by a named entity recognition system (as we explain in the next section). The following attempts are then made for the verb-relation mapping:

1) Search the KB for an exact matching of the verbal expression with any existent relation for the instances under consideration or any possible relation between the classes (and superclasses) of the instances under consideration.

2) Apply the string similarity metrics to calculate the similarity between the given verbal expression and the possible relations between instances (or their classes) corresponding to the terms in the linguistic triple.

3) Search for similar mappings for the types/classes of entities under consideration in a lexicon of mappings previously accomplished according to users' choices in the question answering system Aqualog. This lexicon contains ontology triples along with the original verbal expression, as illustrated in Table 1. The use of this lexicon represents a simplified form of pattern matching in which only exact matching is considered.

given_relation	class_1	conceptual_relation	class_2
works	project	has-project-member	person
cite	project	has-publication	publication

Table 1. Examples of lexicon patterns

4) Search for synonyms of the given verbal expression in WordNet, in order to verify if there is a synonym that matches (complete or partially, using string similarity metrics) any existent relation for the instances under consideration, or any possible relation between the classes (or superclasses) of those instances (likewise in step 1).

If there is no possible mapping for the term, the pattern-based classification model is triggered (Section 3.4). Conversely, if there is more than one possible mapping, the disambiguation model is called (Section 3.5).

The application of these strategies to map the linguistic triples into existent or new instances and relations is described in what follows.

Applying RSS to map entities and relations

In our architecture, RSS is represented by modules **RSS_1** and **RSS_2**. **RSS_1** first checks if

the terms in the linguistic triple are instances of a KB (cf. strategies for mapping terms). If the terms can be mapped to instances, it checks whether the relation given in the triple matches any already existent relation between for those instances, or, alternatively, if that relation matches any of the possible relations for the classes (and superclasses) of the two instances in the domain ontology (cf. cf. strategies for mapping relations). Three situations may arise from this attempt to map the linguistic triple into an ontology triple (Cases (1), (2), and (3) in Fig. 1):

Case (1): complete matching with instances of the KB and a relation of the KB or ontology, with possibly more than one valid conceptual relation being identified:

$\langle \text{instance}_1, (\text{conceptual_relation})^+, \text{instance}_2 \rangle$.

Case (2): no matching or partial matching with instances of the ontology (the relation is not analyzed (*na*) when there is not a matching for instances):

$\langle \text{instance}_1, na, ? \rangle$ or $\langle ?, na, \text{instance}_2 \rangle$ or $\langle ?, na, ? \rangle$

Case (3): matching with instances of the KB, but no matching with a relation of the KB or ontology:

$\langle \text{instance}_1, ?, \text{instance}_2 \rangle$

If the matching attempt results in Case (1) with only one conceptual relation, then the triple can be formalized into a semantic annotation. This yields the annotation of an already existent relation for two instances of the KB, as well as a new relation for two instances of the KB, although this relation was already predicted in the ontology as possible between the classes of those instances. The generalization of the produced triple for classes/types of entities, i.e., $\langle \text{class}, \text{conceptual_relation}, \text{class} \rangle$, is added to the repository of **Patterns**.

On the other hand, if there is more than one possible conceptual relation in case (1), the system tries to find the correct one by means of a sense disambiguation model, described in Section 3.5. Conversely, if there is no matching for the relation (Case (3)), the system tries an alternative strategy: the pattern-based classification model (Section 3.4). Finally, if there is no complete matching of the terms with instances of the KB (Case (2)), it means that the entities can be new to the KB.

In order to check if the terms in the linguistic triple express new entities, the system first iden-

tifies to what classes of the ontology they belong. This is accomplished by means of ESpotter++, and extension of the named entity recognition system ESpotter (Zhu et al, 2005).

ESpotter is based on a mixture of lexicon (gazetteers) and patterns. We extended ESpotter by including new entities (extracted from other gazetteers), a few relevant new types of entities, and a small set of efficient patterns. All types of entities correspond to generic classes of our domain ontology, including: person, organization, event, publication, location, project, research-area, technology, etc.

In our architecture, if ESpotter++ is not able to identify the types of the entities, the process is aborted and no annotation is produced. This may be either because the terms do not have any conceptual mapping (for example “it”), or because the conceptual mapping is not part of our domain ontology. Otherwise, if ESpotter++ succeeds, RSS is triggered again (**RSS_2**) in order to verify whether the verbal expression encompasses a semantic relation. Since at least one of the two entities is recognized by Espotter++, and therefore at least one entity is new, it is only possible to check if the relation matches the possible relations between the classes of the recognized entities (cf. strategies for mapping relations).

If the matching attempt results in only one conceptual relation, then the triple will be formalized into a semantic annotation. This represents the annotation of a new (although predicted) relation and two or at least one new entity/instance. The produced triple of the type <class, conceptual_relation, class> is added to the repository of **Patterns**.

Again, if there are multiple valid conceptual relations, the system tries to find the correct one by means of a disambiguation model (Section 3.5). Conversely, if there is no matching for the relation, the pattern-based classification model is triggered (Section 3.4).

3.4 Identifying new relations

The process described in Section 3.3 for the identification of relations accounts only for the relations already predicted as possible in the domain ontology. However, we are also interested in the additional information that can be provided by the text, in the form of new types of relations for known or new entities. In order to discover these relations, we employ a pattern matching strategy to identify relevant relations between types of terms.

The pattern matching strategy has proved to be an efficient way to extract semantic relations, but in general has the drawback of requiring the possible relations to be previously defined. In order to overcome this limitation, we employ a **Pattern-based classification** model that can identify similar patterns based on a very small initial number of patterns.

We consider patterns of relations between types of entities, instead of the entities themselves, since we believe that it would be impossible to accurately judge the similarity for the kinds of entities we are addressing (names of people, locations, etc). Thus, our patterns consist of triples of the type <class, conceptual_relation, class>, which are compared against a given triple using its classes (already provided by the linguistic component or by ESpotter++) in order to classify relations in that triple as *relevant* or *non-relevant*.

The classification model is based on the approach presented in (Stevenson, 2004). It is an unsupervised corpus-based module which takes as examples a small set of relevant SVO patterns, called seed patterns, and uses a WordNet-based semantic similarity measure to compare the pattern to be classified against the relevant ones. Our initial seed patterns (see examples in Table 2) mixes patterns extracted from the lexicon generated by Aqualog’s users (cf. Section 3.3) and a small number of manually defined relevant patterns. This set of patterns is expected to be enriched with new patterns as our system annotates relevant relations, since the system adds new triples to the initial set of patterns.

class 1	conceptual relation	class 2
project	has-project-member	person
project	has-publication	publication
person	develop	technology
person	attend	event

Table 2. Examples of seed patterns

Likewise (Stevenson, 2004), we use a semantic similarity metric based on the information content of the words in WordNet hierarchy, derived from corpus probabilities. It scores the similarity between two patterns by computing the similarity for each pair of words in those patterns. A threshold of 0.90 for this score was used here to classify two patterns as similar. In that case, a new annotation is produced for the input triple and it is added to the set of patterns.

It is important to notice that, although WordNet is also used in the RSS module, in that case

only synonyms are checked, while here the similarity metric explores deeper information in WordNet, considering the meaning (senses) of the words. It is also important to distinguish the semantic similarity metrics employed here from the string metrics used in RSS. String similarity metrics simply try to capture minor variations on the strings representing terms/relations, they do not account for the meaning of those strings.

3.5 Disambiguating relations

The ambiguity arising when more than one possible relation exists for a pair of entities is a problem neglected in most of the current work on relation extraction. In our architecture, when the RSS finds more than one possible relation, we choose one relation by using the word sense disambiguation (WSD) system SenseLearner (Mihalcea and Csomai, 2005).

SenseLearner is supervised WSD system to disambiguate all open class words in any given text, after being trained on a small data set, according to global models for word categories. The current distribution includes two default models for verbs, which were trained on a corpus containing 200,000 content words of journalistic texts tagged with their WordNet senses. Since SenseLearner requires a sense tagged corpus in order to be trained to specific domains and there is not such a corpus for our domain, we use one of the default training models. This is a contextual model that relies on the first word before and after the verb, and its POS tags. To disambiguate new cases, it requires only that the words are annotated with POS tags. The use of lemmas of the words instead of the words yields better results, since the models were generated for lemmas. In our architecture, these annotations are produced by the component **POS + Lemmatizer**.

Since the WSD module disambiguates among WordNet senses, it is employed only after the use of the WordNet subcomponent by RSS. This subcomponent finds all the synonyms for the verb in a linguistic triple and checks which of them matches existent or possible relations for the terms in that triple. In some cases, however, there is a matching for more than one synonym. Since in WordNet synonyms usually represent different uses of the verb, the WSD module can identify in which sense the verb is being used in the sentence, allowing the system to choose one among all the matching options.

For example, given the linguistic triple <enrico_motta, head, kmi>, RSS is able to identify that “enrico_motta” is a *person*, and that “kmi” is

an *organization*. However, it cannot find an exact or partial matching (using string metrics), or even a matching (given by the user lexicon) for the relation “head”. After getting all its synonyms in WordNet, RSS verifies that two of them match possible relations in the ontology between a *person* and an *organization*: “direct” and “lead”. In this case, the WSD module disambiguates the sense of “head” as “direct”.

3.6 Example of extracted relations

As an example of the relations that can be extracted in our approach, consider the representation of the entity “Enrico Motta” and all the relations involving this entity in Figure 5. The relations were extracted from the text in Figure 6.

KMi awarded £4M for Semantic Web Research

Professor Enrico Motta and Dr John Domingue of the Knowledge Media Institute have received a set of record-breaking awards totalling £4m from the European Commission's Framework 6 Information Society Technologies (IST) programme. This is the largest ever combined award obtained by KMi associated with a single funding programme. The awards include three Integrated Projects (IPs) and three Specific Targeted Research Projects (STREPs) and they consolidate KMi's position as one of the leading international research centers in semantic technologies. Specifically Professor Motta has been awarded:

- a.. £1.55M for the project NeOn: Lifecycle Support for Networked Ontologies
- b.. £565K for XMEDIA: Knowledge Sharing and Reuse across Media and
- c.. £391K for OK: Openknowledge - Open, coordinated knowledge sharing architecture. ...

Figure 5. Example of newsletter

```
(def-instance Enrico-Motta kmi-academic-staff-member
  ((works-in knowledge-media-institute)
   (award-from european-commission)
   (award-for NeOn)
   (award-for XMEDIA)
   (award-for OK)))
```

Figure 6. Semantic annotations produced for the news in Figure 5

In this case, “Enrico-Motta” is an instance of *kmi-academic-staff-member*, a subclass of *person* in the domain ontology. The mapped relation “works-in” “knowledge-media-institute” already existed in the KB. The new relations pointed out by our approach are the ones referring to the award received from the “European Commission” (an *organization*, here), for three *projects*: “NeOn”, “XMEDIA”, and “OK”.

4 Conclusions and future work

We presented a hybrid approach for the extraction of semantic relations from text. It was de-

signed mainly to enrich the annotations produced by a semantic web portal, but can be used for other domains and applications, such as ontology population and development. Currently we are concluding the integration of the several modules composing our architecture. We will then carry experiments with our corpus of newsletters in order to evaluate the approach. Subsequently, we will incorporate the architecture to a semantic web portal and accomplish an extrinsic evaluation in the context of that application. Since the approach uses deep linguistic processing and corpus-based strategies not requiring any manual annotation, we expect it will accurately discover most of the relevant relations in the text.

Acknowledgement

This research was supported by the Advanced Knowledge Technologies (AKT) project. AKT is an Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01.

References

- Massimiliano Ciaramita, Aldo Gangemi, Esther Ratsch, Jasmim Saric, Isabel Rojas. 2005. Unsupervised learning of semantic relations between concepts of a molecular biology ontology. *19th IJCAI*, pp. 659-664
- Hamish Cunningham, Diana Maynard, Kalina Bontcheva, and Valentin Tablan. 2002. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. *40th ACL Meeting*, Philadelphia.
- Hamish Cunningham, Diana Maynard, and Valentin Tablan. 2000. JAPE: a Java Annotation Patterns Engine. *Tech. Report CS--00--10*, University of Sheffield, Department of Computer Science.
- Christiane D. Fellbaum (ed). 1998. *Wordnet: An Electronic Lexical Database*. The MIT Press.
- Pablo Gamallo, Marco Gonzalez, Alexandre Agustini, Gabriel Lopes, and Vera S. de Lima. 2002. Mapping syntactic dependencies onto semantic relations. *ECAI Workshop on Machine Learning and Natural Language Processing for Ontology Engineering*, Lyon, France.
- Asuncion Gomez-Perez and David Manzano-Macho. 2003. *A Survey of Ontology Learning Methods and Techniques*. Deliverable 1.5, OntoWeb Project.
- Jose Iria and Fabio Ciravegna. 2005. Relation Extraction for Mining the Semantic Web. *Dagstuhl Seminar on Machine Learning for the Semantic Web*, Dagstuhl, Germany.
- Yuanguai Lei, Marta Sabou, Vanessa Lopez, Jianhan Zhu, Victoria Uren, and Enrico Motta. 2006. An infrastructure for Acquiring High Quality Semantic Metadata. To appear in the *3rd ESWC*, Budva.
- Dekang Lin. 1993. Principle based parsing without overgeneration. *31st ACL*, Columbus, pp. 112-120.
- Vanessa Lopez, Michele Pasin, and Enrico Motta. 2005. AquaLog: An Ontology-portable Question Answering System for the Semantic Web. *2nd ESWC*, Crete, Grece.
- Alexander D. Maedche. 2002. *Ontology Learning for the Semantic Web*, Kluwer Academic Publishers, Norwell, MA.
- Scott Miller, Heidi Fox, Lance Ramshaw, and Ralph Weischedel. 2000. A novel use of statistical parsing to extract information from text. *6th ANLP-NAACL*, Seattle, pp. 226-233.
- Rada Mihalcea and Andras Csomai. 2005. Sense-Learner: Word Sense Disambiguation for All Words in Unrestricted Text. *43rd ACL Meeting*, Ann Arbor.
- Marie-Laure Reinberger and Peter Spyns. 2004. Discovering knowledge in texts for the learning of DOGMA inspired ontologies. *ECAI 2004 Workshop on Ontology Learning and Population*, Valencia, pp. 19-24.
- Dan Roth and Wen-tau Yih. 2002. Probabilistic reasoning for entity & relation recognition. *19th COLING*, Taipei, Taiwan, pp. 1-7.
- Alexander Schutz and Paul Buitelaar. 2005. RelExt: A Tool for Relation Extraction from Text in Ontology Extension. *4th ISWC*, pp. 593-606.
- Mark Stevenson. 2004. An Unsupervised WordNet-based Algorithm for Relation Extraction. *4th LREC Workshop Beyond Named Entity: Semantic Labeling for NLP Tasks*, Lisbon.
- Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. 2003. Kernel Methods for Relation Extraction. *Journal of Machine Learning Research*, (3):1083-1106.
- Shubin Zhao and Ralph Grishman. 2005. Extracting Relations with Integrated Information Using Kernel Methods. *43d ACL Meeting*, Ann Arbor.
- Jianhan Zhu, Victoria Uren, and Enrico Motta. 2005. ESpotter: Adaptive Named Entity Recognition for Web Browsing. *3rd Conf. on Professional Knowledge Management*, Kaiserslautern, pp. 518-529.
- Roman Yangarber, Ralph Grishman and Pasi Tapanainen, P. 2000. Unsupervised Discovery of Scenario-Level Patterns for Information Extraction. *6th ANLP*, pp. 282-289.