IXAGroupEHUSpaceEval: (*X-Space*) A WordNet-based approach towards the Automatic Recognition of Spatial Information following the ISO-Space Annotation Scheme

Haritz Salaberri, Olatz Arregi, Beñat Zapirain
IXA Group, Faculty of Computer Sciences
University of the Basque Country (UPV-EHU)
Manuel Lardizabal pasealekua, 1 - 20018 Donostia-San Sebastián
`{haritz.salaverri, olatz.arregi, benat.zapirain}@ehu.eus`

Abstract

This paper presents *X-Space*, a system that follows the ISO-Space annotation scheme in order to capture spatial information as well as our contribution to the SemEval-2015 task 8 (SpaceEval). Our system is the only participant system that reported results for all three evaluation configurations in SpaceEval.

1 Introduction

Nowadays the need for algorithms that have the ability to reason spatially over texts are in growing demand within applications concerning human language processing and in navigation services. Well-studied topics in computational linguistics such as named entity recognition and question answering, for example, will presumably experience important progress through such algorithms. Navigation systems, on the other hand, will gain the ability to interpret indications given by users beyond the "string matching" methods used at present (Wu et al., 2010). In order for such systems to reason spatially, however, they require the enrichment of textual data with the annotation of spatial information in language (Pustejovsky et al., 2013). As of today there have been several attempts at capturing spatial information (annotation schemes): SpatialML (Mani et al., 2008), Spatial Role Labeling (Kordjamshidi et al., 2010) and ISO-Space (Pustejovsky et al., 2011). *X-Space* follows the ISO-Space specification.

Section 2 describes the system architecture and section 3 presents the results and discusses them. Finally, our conclusions are given in section 4.

2 System architecture

X-Space uses a four-stage pipeline: first texts are preprocessed; second candidates to be spatial elements and signals are selected by generating word lists from the texts; then, spatial elements and signals are identified from the candidate lists and their attributes set according to type. Finally, spatial relations are established between the previously identified and the attributes that correspond to these relations are set.

2.1 Preprocessing and Candidate Selection

As an initial step texts that are inputted into *X-Space* are syntactically and semantically parsed (SRL), and named entities as well as coreference chains are identified. These annotations are used in the later stages as features for machine learning.

The parsing of syntactic and semantic dependencies is achieved with the *ClearNLP* semantic role labeler (Choi and Palmer, 2012); for the recognition of named entities, on the other hand, the *Apache OpenNLP* name finder tool is used (Baldridge, 2005). Chains of coreference are identified using the *Standford CoreNLP* coreference resolution system (Manning et al., 2014).

After the preprocessing of input texts is performed the words that compose these texts are used to form candidate lists by taking words one-by-one, two-by-two, three-by-three and four-by-four. We assume that spatial elements and signals with more than four words are highly improbable to occur. For this reason only candidates with up to four words are considered.

856

2.2 Spatial Elements and Signals

Five different spatial elements are distinguished: places and paths which designate a region of space (locations); spatial entities, words of motion and non-motion events. These do not designate a region of space but are allowed to be coerced into behaving like a region of space, so that they may participate in the same kinds of relationships as regions of space (Pustejovsky and Yocum, 2013).

In order to identify places, paths and words of motion we have used WordNet as well as several other resources such as PropBank and the Predicate Matrix (de Lacalle et al., 2014) in combination with a binary *Support Vector Machine classifier* implemented using the *SVM-light* package (Joachims, 1999). For spatial entities and non-motion events, on the other hand, an approach without WordNet is used for reasons that are discussed in 2.2.2.

2.2.1 The WordNet approach

The WordNet approach is used to identify words of motion, paths and places within the lists of candidates. This approach is based on the idea that within the hierarchical organization of WordNet a domain, as for example the path domain, can be defined as a set of subtrees by properly identifying the root of these subtrees. According to (Feizabadi and Padó, 2012) the challenge is to find a set of nodes whose subtrees cover as much as possible of the desired domain while avoiding overgeneration.

Root nodes We consider these nodes to be the ones that best fulfill these conditions after manually examining WordNet (v3.0): for the motion domain we have considered nodes "move, locomote, travel, go" (01835496-V) and "to be" (02604760-V); for the place domain, on the other hand, we have considered nodes "topographic point, place spot" (08664443-N), "place, property" (08513718-N), "position, place" (08621598-N), "location" (00027167-N), "state, nation, country, land, commonwealth, *res publica*, body politic" (08168978-N), "country, state, land" (08544813-N), "country, rural area" (08644722-N) and "area, country" (08497294-N). Finally, for the path domain we have considered nodes "path" (03899328-N), "path, route, itinerary" (08616311-N), "path, track, course" (09387222-N) and "way" (04564698-N).

Domain definition After the place and path domains are defined by capturing the corresponding sets of subtrees, the domains are completed by adding the places and paths in the training set that are not covered by the subtrees. In total 5,572 places and 664 paths form the final domains. For motion words, however, most of them being verbs, sense needs to be disambiguated; as a matter of fact many verbs have motion as well as non-motion senses. *X-Space* uses the Predicate Matrix (v1.1) in order to map the WordNet synset IDs that correspond to the subtrees rooted in the nodes that have been considered for the motion domain (01835496-V and 02604760-V) with their corresponding PropBank sense. The Predicate Matrix is a lexical resource resulting from the integration of multiple sources of predicate information, including FrameNet, VerbNet, PropBank and WordNet. 511 senses form the words of motion domain used by our system.

Identification process When domains are defined *X-Space* iterates over the list of candidates in search for places and paths. Words of motion, on the contrary, are only looked for in the one-by-one candidate list as we consider that only these can be words of motion. For the identification of words of motion the disambiguation of senses is used which is performed in the semantic dependency parsing in the preprocessing stage. This way only the candidates labeled with a sense present in the words of motion domain will be identified as such.

Avoiding overgeneration We observed that too many candidates were identified using just the straightforward procedure, where candidates are looked for in the domains. In order to avoid this over-generation we attached a machine-learning module to the identification process. The classifier used is a binary classifier based on *Support Vector Machines*, and once the identification of places, paths and words of motion within the identification process is completed, the classifier decides whether an identified element is or not correctly identified. The following is the list of the features used:

- CAND_Lex, CAND_Lemma, CAND_PoS: Lexical form, lemma and PoS category tag of

the candidate.

- CAND_DepRel: Dependency relation between the candidate and its head.

- PRED_Roleset: Roleset ID of the predicate on which the candidate semantically depends.

- PREVW_Lex, PREVW_PoS: Lexical form and PoS category tag of the word previous to the candidate.

- PREVW_DepRel: Dependency relation between the word previous to the candidate and its head.

- NEXTW_Lex, NEXTW_PoS: Lexical form and PoS category tag of the word next to the candidate.

- NEXTW_DepRel: Dependency relation between the word next to the candidate and its head.

X-Space uses this classifier (same set of features, same implementation) several times throughout the entire process of annotating input texts with spatial information. The reason not to perform a thorough feature selection process whenever machine learning needs to be used throughout the annotation conducted by *X-Space* lies in the time limitations we encountered due to the extent of the SpaceEval task. The values taken by the training parameters are shown next:

- Trade-Off: The trade-off between training error and margin is computed through the $avg(x * x)^{-1}$ formula.

- Bias: A biased hyperplane is used.

- Cost-Factor: The cost-factor, by which training errors on positive examples outweigh errors on negative examples is 1.

- Kernel: The type of kernel function used is linear.

2.2.2 Other approaches

This section describes the approaches used to identify non-motion events, spatial entities and signals from the candidates lists. The reason why the WordNet approach is not used on non-motion events is that we could not accurately determine a set of nodes whose subtrees cover the desired domain properly. For spatial entities, on the other hand, we believe that using the WordNet approach is not correct given the heterogeneous nature of these spatial elements. We also believe that the same reason applies not to use WordNet in the identification of spatial and motion signals. All three use a classifier like the one in section 2.2.1 as a final step to avoid overgeneration.

Non-motion events   In order to identify non-motion events our system first generates a list with the PropBank senses taken by the non-motion events in the training set. Then *X-Space* iterates over the one-by-one candidate list (as we consider that only these can be ISO-Space non-motion events) and checks whether a candidate is labeled with one of these senses or not. With this aim in mind the disambiguation of senses is used that has been performed in the preprocessing.

Spatial entities   For the identification of spatial entities the semantic role labels have been used that were given by the semantic dependency parser to predicate arguments. We believe spatial entities to be viewable as arguments of a predicate which is a word of motion or a word that expresses a non-motion event. Arguments that correspond to these kind of predicates are in the majority of cases located in space and participate in ISO-Space link tags. These, therefore, can be understood as spatial entities. In order to identify spatial entities *X-Space* iterates over the lists of candidates and searches for arguments of words that have been previously marked as of motion or as expressing a non-motion event.

Signals   For the purpose of identifying signals two lists are formed based on the signal annotations in the training set. One list holds the signals that are exclusively of motion (e.g., into,

from) and another list holds the signals that can only be of space (e.g., within, without). Then *X-Space* iterates over the one-by-one and two-by-two candidate lists. As we observed in the training set, only signals with up to two words occurred. In the iteration process candidates are looked for in both lists; if a candidate can be found in the spatial signals list, it is assigned the spatial signal tag; on the other hand, if the candidate is present in the motion signals list this is marked as a motion signal. Many signals, however, overlap, meaning that they can be of motion and space (e.g., by, over). In order to capture these signals the candidate lists are searched for prepositions and function words.

2.2.3 Attribute identification

The ISO-Space annotation language specifies several attributes for spatial elements as well as for signals. The classifier in section 2.2.1 is used to give values to these attributes. When an attribute can take more than two values, however, a version of the classifier that has been extended from binary to multiclass is used. This is achieved with the *SVM-multiclass* package. Several attributes are specified. However, not all attributes are given values: in fact many of them are never annotated in the training set, and for this reason *X-Space* does not annotate them either.

2.3 Spatial Relation Links

When the spatial elements and the signals within the input texts as well as their corresponding attributes are identified *X-Space* tries to detect the spatial relations that lay between them. The SpaceEval task addresses the detection of three types of relations: movement (MoveLink), qualitative (QSLink) and orientational (OLink).

2.3.1 Link identification

In order to identify the spatial relation links, *X-Space* follows what is stated in (Pustejovsky and Yocum, 2013). For movement relations that typically involve motion-event triggers (words of motion), motion signals, and motion-event participants a link is created for each identified word of motion. For qualitative relations, on the other hand, which normally involve spatial signals and spatial elements and are used to capture topological relationships,

a link of this type is created for each spatial signal with an identified TOPOLOGICAL or DIR_TOP `semantic_type`. Finally, an orientational link is introduced for every spatial signal with a DIRECTIONAL `semantic_type`; this kind of link describes non-topological relationships between spatial signals and spatial elements.

2.3.2 Attribute identification

There are several attributes specified by the ISO-Space annotation scheme for the MoveLink, QSLink and OLink relations. Nonetheless, not all these attributes are viewed, and, consequently, identified by *X-Space* following the same procedure. In fact, the system distinguishes three types of attributes: triggers, which are the spatial elements or signals that trigger the creation of links, roles, which are the spatial elements involved in these relations and common attributes, which indicate other characteristics of the links.

Triggers Triggers are directly established for all three kinds of links based on the link identification process.

Roles We believe that attributes `source`, `goal`, `mover` and `landmark` within a MoveLink relation can be seen as arguments of the trigger, which is usually a verbal predicate (word of motion). For a QSLink or OLink relation, on the other hand, we think that attributes `trajector` and `landmark` can be seen as arguments of the predicate that dominates the trigger, which is usually a preposition (spatial signal). The idea behind these attributes is based on the Spatial Role Labeling annotation schema described in (Kordjamshidi et al., 2010). According to this scheme there are indicators that can be spatial (spatial signals) or of motion (words of motion) that introduce spatial relations. These spatial relations take arguments with roles `trajector` and `landmark`. For this reason, in order to identify attributes `source`, `goal`, `mover` and `landmark` of MoveLinks, *X-Space* looks for arguments of the triggers using the semantic dependency parsing carried out in the preprocessing. Then it establishes which PropBank argument (A0, A1, A2, etc.) corresponds to which attribute (spatial role) using a multiclass classifier

	Precision		Recall		F_1		Accuracy	
	Baseline	X-Space	Baseline	X-Space	Baseline	X-Space	Baseline	X-Space
1-a	0.55	0.81	0.52	0.72	0.53	0.76	0.75	0.88
1-b	0.55	0.75	0.51	0.72	0.53	0.74	0.86	0.9
1-c	0.1	0.18	0.02	0.15	0.04	0.16	0.05	0.3
1-d	0.5	0.54	0.5	0.51	0.5	0.53	0.5	0.55
1-e	0.05	0.06	0.02	0.05	0.02	0.05	0.06	0.25
2-a	0.27	0.26	0.28	0.33	0.27	0.29	0.76	0.63
2-b	0.79	0.55	0.58	0.51	0.67	0.53	0.9	0.89
2-c	0.19	0.06	0.2	0.08	0.19	0.07	0.66	0.46
3-a	0.86	0.63	0.84	0.51	0.85	0.56	0.98	0.89
3-b	0.26	0.07	0.26	0.09	0.26	0.08	0.79	0.48

Table 1: Official results reported for *X-Space* plus the results of one baseline system provided by the task organizers (overall results).

like the one in section 2.2.1. The procedure for QSLinks and OLinks is the same but arguments are searched for the predicate that dominates the trigger and not for the trigger itself.

Common attributes   We have named common attributes all attributes taken by spatial relations that do not indicate a spatial role or a trigger. *X-Space* once again uses the classifier on 2.2.1 in order to give values to these attributes.

3   Results

The SpaceEval task considers three separate evaluation configurations: (1) only unannotated text is given as an input; (2) manually annotated spatial element extents (no attributes) are given; (3) manually annotated spatial element extents and their attributes are given.

The subtasks that are evaluated for each configuration are: (1-a) identifying spans of spatial elements, (1-b) classifying spatial elements according to type, (1-c) identifying attributes for spatial elements according to type, (1-d) identifying MoveLink, QSLink and OLink relations and (1-e) identifying attributes for spatial relations; (2-a) classifying spatial elements and identifying their attributes according to type, (2-b) identifying MoveLink, QSLink and OLink relations and (2-c) identifying attributes for spatial relations; (3-a) identifying MoveLink, QSLink and OLink relations and (3-b) identifying attributes for spatial relations.

Table 1 shows the results obtained by *X-Space* in every configuration and subtask and compares them with the results of one baseline system provided by the task organizers. As can be noted our results improve the ones in the baseline for 1-a, 1-b, 1-c, 1-d, 1-e and 2-a. On the other hand, our results are worse for 2-b, 2-c, 3-a and 3-b. From the three systems that participated in the SpaceEval task ours was the only one that presented results for all evaluation configurations and all subtasks. We believe that in general the results for our system are good.

4   Conclusion and Future Works

In this paper *X-Space*, our contribution to the SemEval-2015 Task 8, is presented. We consider that many things still remain to be improved. For instance, the problem of annotating non-consuming location tags could be addressed.

In the future, we intend to adapt our system to other languages. This adaptation will bring the opportunity to see how *X-Space* adapts to languages of different natures.

Acknowledgments

References

Jason Baldridge. 2005. The opennlp project. *URL: http://opennlp. apache. org/index. html,(accessed 2 February 2012).*

Jinho D. Choi and Martha Palmer. 2012. Optimization of natural language processing components for robustness and scalability.

Maialen. L. de Lacalle, Egoitz Laparra, and German Rigau. 2014. Predicate matrix: extending semlink through wordnet mappings. In *Proceedings of the 9th conference on International Language Resources and Evaluation (LREC'14).*

Parvin Sadat Feizabadi and Sebastian Padó. 2012. Automatic identification of motion verbs in wordnet and framenet. In *Empirical Methods in Natural Language Processing*, page 70.

Thorsten Joachims. 1999. Making large scale svm learning practical.

Parisa Kordjamshidi, Marie-Francine Moens and Martijn van Otterlo. 2010. Spatial role labeling: Task definition and annotation scheme. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, pages 413–420.

Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby and Ben Wellner. 2008. Spatialml: Annotation scheme, corpora, and tools. In *LREC*.

Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.

James Pustejovsky, Jessica Moszkowicz and Marc Verhagen. 2013. A linguistically grounded annotation language for spatial information.

James Pustejovsky, Jessica Moszkowicz and Marc Verhagen. 2011. Iso-space: The annotation of spatial information in language. In *Proceedings of the Sixth Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation*, pages 1–9.

James Pustejovsky and Zachary Yocum. 2013. Capturing motion in iso-spacebank. In *Workshop on Interoperable Semantic Annotation*, page 25.

Yunhui Wu, Stephan Winter, John A. Bateman, Anthony G. Cohn and James Pustejovsky. 2010. Interpreting place descriptions for navigation services. In *Dagstuhl Seminar on Spatial Representation and Reasoning in Language: Ontologies and Logics of Space, Schloss Dagstuhl, Germany.*