

Elhuyar-IXA: Semantic Relatedness and Cross-Lingual Passage Retrieval

Eneko Agirre¹, Olatz Ansa¹, Xabier Arregi¹, Maddalen Lopez de Lacalle²,
Arantxa Otegi¹, Xabier Saralegi², and Hugo Zaragoza³

¹ IXA NLP Group, University of the Basque Country. Donostia, Basque Country
{e.agirre, olatz.ansa, xabier.arregi, arantza.otegi}@ehu.es

² R&D, Elhuyar Foundation. Usurbil, Basque Country
{maddalen, xabiers}@elhuyar.com

³ Yahoo! Research. Barcelona, Spain
hugoz@yahoo-inc.com

Abstract. This article describes the participation of the joint Elhuyar-IXA group in the ResPubliQA exercise at QA&CLEF. In particular, we participated in the English–English monolingual task and in the Basque–English cross-lingual one. Our focus has been threefold: (1) to check to what extent information retrieval (IR) can achieve good results in passage retrieval without question analysis and answer validation, (2) to check Machine Readable Dictionary (MRD) techniques for the Basque to English retrieval when faced with the lack of parallel corpora for Basque in this domain, and (3) to check the contribution of semantic relatedness based on WordNet to expand the passages to related words. Our results show that IR provides good results in the monolingual task, that our crosslingual system performs lower than the monolingual runs, and that semantic relatedness improves the results in both tasks (by 6 and 2 points, respectively).

1 Introduction

The joint team was formed by two different groups, on the one hand the Elhuyar Foundation, and on the other hand the IXA NLP group. This collaboration allowed us to tackle the English–English monolingual task and the Basque–English cross-lingual one in the ResPubliQA track.

With respect to the Basque-English task, we met the challenge of retrieving English passages for Basque questions. We tackled this problem by translating the lexical units of the questions into English. The main setback is that no parallel corpus was available for this pair of languages, given that there is no Basque version of the JRC-Acquis collection. So we have explored an approach which does not use parallel corpora when translating queries, which could also be interesting for other less resourced languages. In our opinion, bearing in mind the idiosyncrasy of the European Union, it is worthwhile dealing with the search of passages that answer questions formulated in unofficial languages.

Question answering systems typically rely on a passage retrieval system. Given that passages are shorter than documents, vocabulary mismatch problems are more important than in full document retrieval. Most of the previous work on expansion techniques has focused on pseudo-relevance feedback and other query expansion techniques. In particular, WordNet has been used previously to expand the terms in the query with little success [2, 3, 4]. The main problem is ambiguity, and the limited context available to disambiguate the word in the query effectively. As an alternative, we felt that passages would provide sufficient context to disambiguate and expand the terms in the passage. In fact, we do not do explicit word sense disambiguation, but rather apply a state-of-the-art semantic relatedness method [5] in order to select the best terms to expand the documents.

2 System Overview

2.1 Question Pre-processing

We analysed the Basque questions by re-using the linguistic processors of the *Ihardetsi* question-answering system [1]. This module uses two general linguistic processors: the lemmatizer/tagger named *Morfeus* [6], and the Named Entity Recognition and Classification (NERC) processor called *Eihera* [7]. The use of the lemmatizer/tagger is particularly suited to Basque, as it is an agglutinative language. It returns only one lemma and one part of speech for each lexical unit, which includes single word terms and multiword terms (MWTs) (those included in the Machine Readable Dictionary (MRD) introduced in the next subsection). The NERC processor, *Eihera*, captures entities such as *person*, *organization* and *location*. The numerical and temporal expressions are captured by the lemmatizer/tagger. The questions thus analyzed are passed to the translation module.

English queries were tokenized without further analysis.

2.2 Translation of the Query Terms (Basque-English Runs)

Once the questions had been linguistically processed, we translated them into English. Due to the scarcity of parallel corpora for a small language or even for big languages in certain domains, we have explored a MRD-based method. These approaches have inherent problems, such as the presence of ambiguous translations and also out-of-vocabulary (OOV) words. To tackle these problems, some techniques have been proposed such as structured query-based techniques [8, 9] and concurrences-based techniques [10, 11]. These approaches have been compared for Basque by obtaining best MAP (Mean Average Precision) results with structured queries [12]. However, structured queries were not supported in the retrieval algorithm used (see Section 2.3), so we adopted a concurrences-based translation selection strategy.

The translation process designed comprises two steps and takes the keywords (Name Entities, MWTs and single words tagged as noun, adjective or verb) of the question as source words.

In the first step the translation candidates of each source word are obtained. The translation candidates for the lemmas of the source words are taken from a bilingual eu-en MRD composed from the Basque-English *Morris* dictionary¹, and the *Euskal-term* terminology bank² which includes 38,184 MWTs. After that, OOV words and ambiguous translations are dealt with. The number of OOV words quantified out of a total of 421 keywords for the 77 questions of the development set was 42 (10%). Nevertheless, it must be said that many of these OOV words were wrongly tagged lemmas and entities. We deal with OOV words by searching for their cognates in the target collection. The cognate detection is done in two phases. Firstly, we apply several transliteration rules to the source word. Then we calculate the Longest Common Subsequence Ratio (LCSR) among words with a similar length (+-10%) from the target collection (see Figure 1). The ones which reach a previously established threshold (0.9) are selected as translation candidates. The MWTs that are not found in the dictionary are translated word by word, as we realized that most of the MWTs could be translated correctly in that way, exactly 91% of the total MWTs identified by hand in the 77 development questions.

err- ---> r- *erradioterapeutiko=radioterapeutiko*
 k ---> c *radioterapeutiko=radioterapeutico*
 $LCSR(radioterapeutico, radioterapeutico) = 0.9375$

Fig. 1. Example of cognate detection

In the second step, we select the best translation of each source keyword according to an algorithm based on target collection concurrences. This algorithm sets out to obtain the translation candidate combination that maximizes their global association degree. We take the algorithm proposed by Monz and Dorr [11].

Initially, all the translation candidates are equally likely. Assuming that t is a translation candidate of the set of all candidates $tr(s_i)$ for a query term s_i given by the MRD, then:

Initialization step:

$$w_T^0(t | s_i) = \frac{1}{|tr(s_i)|} \quad (1)$$

In the iteration step, each translation candidate is iteratively updated using the weights of the rest of the candidates and the weight of the link connecting them.

Iteration step:

$$w_T^n(t | s_i) = w_T^{n-1}(t | s_i) + \sum_{t' \in \text{inlink}(t)} w_L(t, t') w_T^{n-1}(t' | s_i) \quad (2)$$

¹ English/Basque dictionary including 67,000 entries and 120,000 senses.

² Terminological dictionary including 100,000 terms in Basque with equivalences in Spanish, French, English and Latin.

where $inlink(t)$ is the set of translation candidates that are linked to t , and $w_L(t, t')$ is the association degree between t and t' on the target passages measured by Log-likelihood ratio. These concurrences were calculated by taking the target passages as window.

After re-computing each term weight they are normalized.

Normalization step:

$$w_T^n(t | s_i) = \frac{w_T^n(t | s_i)}{\left[\frac{1}{|r(s_i)|} \sum_{m=1}^n w_T^n(t_{i,m} | s_i) \right]} \tag{3}$$

The iteration stops when the variations of the term weights become smaller than a predefined threshold.

We have modified the iteration step by adding a factor $w_F(t, t')$ to increase the association degree $w_T(t, t')$ between translation candidates t and t' whose corresponding source words $so(t), so(t')$ are close to each other (distance dis in words is low) in the source query Q , or even belong to the same Multi-Word Unit ($smw(so(t), so(t'))=1$). As the global association degree between translation candidates is estimated from the association degree of pairs of candidates, we score positively these two characteristics when the association degree for a pair of candidates is calculated. Thus, the modified association degree $w'_L(t, t')$ between t and t' will be calculated in this way:

$$w'_L(t, t') = w_L(t, t') \cdot w_F(t, t') \tag{4}$$

$$w_F(t, t') = \frac{\max_{s_i, s_j \in Q} dis(s_i, s_j)}{dis(so(t), so(t'))} \cdot 2^{smw(so(t), so(t'))} \tag{5}$$

$$smw(s, s') = \begin{cases} 1 & \{s, s'\} \subseteq Z \text{ where } Z \in MWU \\ 0 & \end{cases} \tag{6}$$

2.3 Passage Retrieval

The purpose of the passage retrieval module is to retrieve passages from the document collection which are likely to contain an answer. The main feature of this module is that the passages are expanded based on their related concepts, as explained in the following sections.

2.3.1 Document Preprocessing and Application of Semantic Relatedness

Given that the system needs to return paragraphs, we first split the document collection into paragraphs. Then we lemmatized and part-of-speech (POS) tagged those passages using the OpenNLP open source software³.

³ <http://opennlp.sourceforge.net/>

After preprocessing the documents, we expanded the passages based on semantic relatedness. To this end, we used UKB⁴, a collection of programs for performing graph-based Word Sense Disambiguation and lexical similarity/relatedness using a pre-existing knowledge base [5], in this case WordNet 3.0.

Given a passage (represented using the lemmas of all nouns, verbs, adjectives and adverbs), UKB returns a vector of scores for concepts in WordNet. Each of these concepts has a score, and the higher the score, the more related the concept is to the given passage. Given the list of related concepts, we took the highest-scoring 100 concepts and expanded them to all variants (words that lexicalize the concepts) in WordNet. An example of a document expansion is shown in Figure 2.

We applied the expansion strategy only to passages which had more than 10 words (half of the passages), for two reasons: the first one is that most of these passages were found not to contain relevant information for the task (e.g. “Article 2”, “Having regard to the proposal from the Commission” or “HAS ADOPTED THIS REGULATION”), and the second is that we thus saved some computation time.

2.3.2 Indexing

We indexed the new expanded documents using the MG4J search-engine [13]. MG4J makes it possible to combine several indices over the same document collection. We created one index for the original words and another one with the variants for the most related 100 concepts. This way, we were able to use the original words only, or alternatively, to also include the expanded words during the retrieval. Porter stemmer was used.

2.3.3 Retrieval

We used the BM25 ranking function with the following parameters: 1.0 for $k1$ and 0.6 for b . We did not tune these parameters. MG4J allows multi-index queries, where one can specify which of the indices one wants to search in, and assign different weights to each index. We conducted different experiments, by using only the index made of original words and also by using the index with the expansion of concepts, giving different weights to the original words and the expanded concepts. The weight of the index which was created using the original words from the passages was 1.00 for all the runs. 1.00 was also the weight of the index that included the expanded words for the monolingual run, but it was 1.78 for the bilingual run. These weights were fixed following a training phase with the English development questions provided by the organization, and after the Basque questions had been translated by hand (as no development Basque data was released). The submitted runs are described in the next section.

3 Description of Runs

We participated in the English-English monolingual task and the Basque-English cross-lingual task. We did not analyze the English queries for the monolingual run,

⁴ The algorithm is publicly available at <http://ixa2.si.ehu.es/ukb/>

and we just removed the stopwords. For the bilingual runs, we first analyzed the questions (see Section 2.1), then we translated the question terms from Basque to English (see Section 2.2), and, finally, we retrieved the relevant passages for the translated query terms (see Section 2.3).

As we were interested in the performance of passage retrieval on its own, we did not carry out any answer validation, and we just chose the first passage returned by the passage retrieval module as the response. We did not leave any question unanswered.

For both tasks, the only difference between the submitted two runs is the use (or not) of the expansion in the passage retrieval module. That is, in the first run (“run 1” in Table 1), during the retrieval we only used the original words that were in the passage. In the second run (“run 2” in Table 1), apart from the original words, we also used the expanded words.

4 Results

Table 1 summarizes the results of our submitted runs, explained in Section 3.

Table 1. Results for submitted runs

submitted runs		#answered correctly	#answered incorrectly	c@1
English - English	run 1	211	289	0.42
	run 2	240	260	0.48
Basque - English	run 1	78	422	0.16
	run 2	91	409	0.18

The results show that the use of the expanded words (run 2) was effective for both tasks, improving the final result by 6 % in the monolingual task.

Figure 2 shows an example of a document expansion which was effective for answering the English question number 32: “*Into which plant may genes be introduced and not raise any doubts about unfavourable consequences for people's health?*”

In the second part of the example we can see some words that we obtained after applying the expansion process explained in Section 2.3.1 to the original passage showed in the example too. As we can see, there are some new words among the expanded words that are not in the original passage, such as *unfavourable* or *consequence*. Those two words were in the question we mentioned before (number 32). That could be why we answered that question correctly when using the expanded words (in run 2), but not when using the original words only.

original passage: *Whereas the Commission, having examined each of the objections raised in the light of Directive 90/220/EEC, the information submitted in the dossier and the opinion of the Scientific Committee on Plants, has reached the conclusion that there is no reason to believe that there will be any adverse effects on human health or the environment from the introduction into maize of the gene coding for phosphinotricine-acetyl-transferase and the truncated gene coding for beta-lactamase;*

some expanded words: *cistron factor gene coding cryptography secret_writing ... acetyl acetyl_group acetyl_radical ethanoyl_group ethanoyl_radical beta_lactamase penicillinase ... ec eec eu europe european_community european_economic_community european_union ... directive directing directional guiding citizens_committee committee environment environs surround surroundings corn ... maize zea_mays health wellness health adverse contrary homo human human_being man adverse inauspicious untoward gamboge ... unfavorable unfavourable ... set_up expostulation objection remonstrance remonstratation dissent protest believe light lightly belief feeling impression notion opinion ... reason reason_out argue jurisprudence law consequence effect event issue outcome result upshot ...*

Fig. 2. Example of a document expansion (doc_id: *jrc31998D0293-en.xml*, p_id: 17)

As expected, the best results were obtained in the monolingual task. With the intention of finding reasons to explain the significant performance drop in the bilingual run, we analyzed manually 100 query translations obtained in the query translation process of the 500 test queries, and detected several types of errors arising from both the question analysis process and from the query translation process. In the question analysis process, some lemmas were not correctly identified by the lemmatizer/tagger, and in other cases some entities were not returned by the lemmatizer/tagger causing us to lose important information for the subsequent translation and retrieval processes. In the query translation process, leaving aside the incorrect translation selections, the words appearing in the source questions were not exactly the ones that figured in many queries that had been correctly translated. In most cases this happened because the English source query word was not a translation candidate in the MRD. If we assume that the answers contain words that appear in the questions and therefore in the passage that we must return, this will negatively affect the final retrieval process.

5 Conclusions

The joint Elhuyar-Ixa team has presented a system which works on passage retrieval alone, without any question analysis and answer validation steps. Our English-English results show that good results can be achieved by means of this simple strategy. We experimented with applying semantic relatedness in order to expand passages prior to indexing, and the results are highly positive, especially for English-English. The performance drop in the Basque-English bilingual runs is significant, and is caused by the accumulation of errors in the analysis and translation of the query mentioned.

Acknowledgments

This work has been supported by KNOW (TIN2006-15049-C03-01), imFUTOURnet (IE08-233) and KYOTO (ICT-2007-211423). Arantxa Otegi's work is funded by a

PhD grant from the Basque Government. Part of this work was done while Arantxa Otegi was visiting Yahoo! Research Barcelona.

References

1. Ansa, O., Arregi, X., Otegi, A., Soraluze, A.: Ihardetsi: A Basque Question Answering System at QA@CLEF 2008. In: Peters, C., Deselaers, T., Ferro, N., Gonzalo, J., Jones, G.J.F., Kurimo, M., Mandl, T., Peñas, A., Petras, V. (eds.) *Evaluating Systems for Multilingual and Multimodal Information Access*. LNCS, vol. 5706, pp. 369–376. Springer, Heidelberg (2009) ISSN 0302-9743 ISBN 978-3-642-04446
2. Kim, S., Seo, H., Rim, H.: Information retrieval using word senses: Root sense tagging approach. In: *Proceedings of SIGIR* (2004)
3. Liu, S., Yu, C., Meng, W.: Word Sense Disambiguation in Queries. In: *Proceedings of the 14th ACM Conference on Information and Knowledge Management, CIKM* (2005)
4. Pérez-Agüera, J.R., Zaragoza, H.: Query Clauses and Term Independence. In: *Evaluating Systems for Multilingual and Multimodal Information Access, 9th Workshop of the Cross-Language Evaluation Forum*. LNCS, pp. 369–376. Springer, Heidelberg (2009) ISSN 0302-9743 ISBN 978-3-642-04446
5. Agirre, E., Soroa, A., Alfonseca, E., Hall, K., Kravalova, J., Pasca, M.: A study on similarity and relatedness using distributional and WordNet-based approaches. In: *Proceedings of the Annual Meeting of the North American Chapter of the Association of Computational Linguistics (NAACL)*, Boulder, USA (2009)
6. Ezeiza, N., Aduriz, I., Alegria, I., Arriola, J.M., Urizar, R.: Combining Stochastic and Rule-Based Methods for Disambiguation in Agglutinative Languages. In: *COLING-ACL*, pp.380–384 (1998)
7. Alegria, I., Arregi, O., Balza, I., Ezeiza, N., Fernandez, I., Urizar, R.: Development of a Named Entity Recognizer for an Agglutinative Language. In: *IJCNLP* (2004)
8. Darwish, K., Oard, D.W.: Probabilistic structured Query Methods. In: *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 338–344 (2003)
9. Pirkola, A.: The effects of query structure and dictionary setups in dictionary-based cross-language information retrieval. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 55–63 (1998)
10. Ballesteros, L., Bruce Croft, W.: Resolving Ambiguity for Cross-language Retrieval. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 64–71 (1998)
11. Monz, C., Dorr, B.J.: Iterative translation disambiguation for cross-language Information Retrieval. In: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 520–527 (2005)
12. Saralegi, X., López de Lacalle, M.: Comparing different approaches to treat Translation Ambiguity in CLIR: Structured Queries v. Target Co-occurrence Based Selection. In: *6th TIR Workshop* (2009)
13. Boldi, P., Vigna, S.: MG4J at TREC 2005. In: Voorhees, E.M., Buckland, L.P. (eds.) *The Fourteenth Text REtrieval Conference (TREC 2005) Proceedings, Number SP 500-266 in Special Publications*. NIST (2005), <http://mg4j.dsi.unimi.it/>