

Rapidly Retargetable Interactive Translingual Retrieval

Gina-Anne Levow, Douglas W. Oard, Philip Resnik and Clara I. Cabezas
Institute for Advanced Computer Studies
University of Maryland, College Park, MD 20742

Abstract

This paper describes a system for rapidly retargetable interactive translingual retrieval. Basic functionality can be achieved for a new document language in a single day, and further improvements require only a relatively modest additional investment. We applied the techniques first to search Chinese collections using English queries, and have successfully added French, German, and Italian document collections. We achieve this capability through separation of language-dependent and language-independent components and through the application of asymmetric techniques that leverage an extensive English retrieval infrastructure.

1 Introduction

Our goal is to produce systems that allow interactive users to present English queries and retrieve documents in languages that they cannot read. In this paper we focus on what we call “rapid retargetability”: extending interactive translingual retrieval functionality for a new document language rapidly with few language-specific resources. Our current system can be retargeted to a new language in one day with only one language-dependent resource: a bilingual term list.¹ Our language-independent architecture consists of two main components:

1. Document translation and indexing
2. Interactive retrieval

We describe each of these components, and then conclude by describing our experience with adding French, German and Italian document collections to a system that was originally developed for Chinese.

2 Document Translation and Indexing

We have adopted a document translation architecture for two reasons. First, we support a single query language (English) but multiple document languages, so indexing English terms

¹For Asian languages we also use a language-specific segmentation system.

simplifies query processing (where interactive response time can be a concern). Second, a document translation architecture simplifies the display of translated documents by decoupling the translation and display processes. Gigabyte collections require machine translation that is orders of magnitude faster than present commercial systems. We accomplish this using term-by-term translation, in which the basic data structure is a simple hash table lookup. Any translation requires some source of translation knowledge—we use a bilingual term list containing English translation(s) for each foreign language term. We typically construct these term lists by harvesting Internet-available translation resources, so the foreign language terms for which translations are known are typically an eclectic mix of root and inflected forms. We accommodate this limitation using a four-stage backoff statistical stemming approach to enhance translation coverage.

Preprocessing. Differences in use of diacritics, case, and punctuation can inhibit matching between term list entries and document terms, so normalization is important. In order to maximize the probability of matching document words with term list entries, we normalize the bilingual term list and the documents by:

- converting characters in Western languages to lowercase,
- removing all accents and diacritics, and
- segmentation, which for Western languages merely requires separating punctuation from other text by the addition of white space.

Our preprocessing also includes conversion of the bilingual term list and the document collection into standard formats. This preprocessing typically requires about half a day of programmer time.

Four-Stage Backoff Translation. Bilingual term lists found on the Web often contain an eclectic mix of root forms and morphological variants. We thus developed a four-stage backoff strategy to maximize coverage while limiting spurious translations:

1. Match the **surface form** of a document term to **surface forms** of source language terms in the bilingual term list.
2. Match the **stem** of a document term to **surface forms** of source language terms in the bilingual term list.
3. Match the **surface form** of a document term to **stems** of source language terms in the bilingual term list.
4. Match the **stem** of a document term to **stems** of source language terms in the bilingual term list.

The process terminates as soon as a match is found at any stage, and the known translations for that match are generated. Although this may produce an inappropriate morphological variant for a correct English translation, use of English stemming at indexing time minimizes

the effect of that factor on retrieval effectiveness. Because we are ultimately interested in processing documents in any language, we may not have a hand-crafted stemmer available for the document language. We have thus explored the application of rule induction to learn stemming rules in an unsupervised fashion from the collection that is being indexed [?].

Balanced Top-2 Translation. We produce exactly two English terms for each foreign-language term. For terms with no known translation, the untranslated term is generated twice (often appropriate for proper names in the Latin-1 character set). For terms with one translation, that translation is generated twice. For terms with two or more known translations, we generate the “best” two translations. In prior experiments we have found that this balanced translation strategy significantly outperforms the usual (unbalanced) technique of including all known translations [?]. We establish the “best” translations by sorting the bilingual term list in advance using only English resources. All single word translations are ordered by decreasing unigram frequency in the Brown corpus, followed by all multi-word translations, and finally by any single word entries not found in the Brown corpus. This ordering has the effect of minimizing the effect of infrequent words in non-standard usages or of misspellings that sometimes appear in bilingual term lists. We translate ~ 200 MB per hour using Perl on a SPARC Ultra 5.

Post-translation Document Expansion. We implement post-translation document expansion for the foreign language stories after translation into English in order to enrich the indexing vocabulary beyond that which was available after term-by-term translation. This is analogous to the process that Singhal et al. applied to monolingual speech retrieval [2].

Term-by-term produces a set of English terms that serve as a noisy representation of the original source language document. These terms are then treated as a query to a comparable English collection, typically contemporaneous newswire text, from which we retrieve the five highest ranked documents. From those five documents, we extract the most selective terms and use them to enrich the original translations of the documents. For this expansion process we select one instance of every term with an IDF value above an *ad hoc* threshold that was tuned to yield approximately 50 new terms. This optional step is the slowest processing stage, with a throughput of about 20 MB per hour. The resulting collection is then indexed using Inquery (version 3.1p1), with the kstem stemmer and default English stopword list. Indexing typically takes well under an hour.

3 Interactive Retrieval

Interactive searches are performed using a Web interface. Summary information (date and gloss translation of the title) for the top-ranked documents is displayed in groups of ten per page. Document summaries consists of the date and a gloss translation of the document title. Users can inspect a gloss translation of the full text of any document if the title is not sufficiently informative. For both title and full text, the gloss translations are generated in advance using the same process as translation for indexing, with the following differences in detail:

- Terms added as a result of document expansion are not displayed.
- The number of retained translations is separately selectable for the title and for full text indexing.
- Translations are not duplicated when fewer than the maximum allowable number of translations are known.

Our goal is to support the process of *finding* documents, with the realization that the process of *using* documents may need to be supported in some other way (e.g., by forwarding relevant documents to someone who is able to read that language). We have therefore designed our interface to highlight the query terms in translated documents and to facilitate skimming by emphasizing the most common translation when multiple translations are displayed. We have found that such displays can support a classification task, even when the translation is not easy to read [1]. Documents must be classified by the user as relevant or not relevant, so our classification results suggest that this is an effective user interface design.

4 Results

We present results both for component-level performance of our language-independent re-targeting modules and an assessment of the overall re-targeting process.

4.1 Component-level Evaluation

We applied our re-targeting approach and retrieval enhancement techniques described above in the context of the first Cross-Language Evaluation Forum’s (CLEF) multilingual task. We used the English language forms of the queries to retrieve English, French, German, and Italian documents. Below we present comparative performance measures for each of the main processing components described above - statistical stemming backoff translation, top-2 balanced translation, and document expansion - applied to the English-French cross-language segment of the CLEF task.

Baseline System Configuration Our baseline run was conducted as follows. We translated the $\sim 44,000$ documents from the 1994 issues of *Le Monde*. We used the English-French bilingual term list downloaded from the Web at <http://www.freedict.com>. We then inverted the term list to form a 35,000 term French-English translation resource. We performed the necessary document and term list normalization; in this case, removing accents from document surface forms to enable matching with the un-accented term list entries, converting case, and splitting clitic contractions, such as *l’horloge*, into its component article and noun or adjective. We trained the statistical stemming rules on a sample of the bilingual term list and document collection and applied these rules in stemming backoff. Our default condition was run with top-2 balanced translation using the Brown corpus as a source of target language unigram frequency information. Translated documents were then indexed with the InQuery (version 3.1p1) system, using the kstem stemmer for English stemming




Figure 1: Comparison of effectiveness of two stemmed versus unstemmed translation of French documents: Bars above x-axis indicate “Stemmed French” outperforms “Unstemmed French”

and InQuery’s default English stopword list. Long queries were formed by concatenating the title, description, and narrative fields of the original query specification. The resulting word sequence was enclosed in an InQuery `#sum` operator, indicating unweighted sum.

Our figure of merit for the evaluations below is mean (uninterpolated) average precision computed using `trec_eval`² across the 34 topics in the CLEF evaluation for which there were relevant French documents.

Statistical Stemming Backoff versus Unstemmed Effectiveness We first contrast the above baseline system with the effectiveness of an otherwise identical run *without* the stemming backoff component. Terms in the documents are thus only translated if there is an exact match between the surface form in the document and a surface form in the bilingual term list. We find that mean average precision for unstemmed translation is 0.19 as compared with 0.2919 for our baseline system including stemming backoff based on trained rules. This difference is significant at $p < 0.05$, by paired t-test, two-tailed. The per-query effectiveness is illustrated in Figure 1 below.

Stemming backoff improves translation coverage while retaining relatively high precision of matching in contrast to unstemmed effectiveness.

Top-1 versus Top-2 Balanced Translation Here we contrast top-2 balanced translation with top-1 translation. We retain statistical stemming backoff for the top-1 translation. We replace each French document term with the highest ranked English translation by target language unigram frequency in the Brown Corpus as detailed above, retaining the original French term when no translation is found in the bilingual term list. We achieve a

²Cornell




Figure 2: Comparison of effectiveness of top-2 balanced versus top-1 translation of French documents: Bars above x-axis indicate “Top-2” outperforms “Top-1”

mean average precision of 0.2532 in contrast with the baseline condition. This difference is significant at $p < 0.01$ by paired t-test, two-tailed. We can effectively incorporate additional translations using top-2 balanced translation without degrading performance by introducing significant additional noise. A query-by-query contrast is presented in Figure reftop1.

Document Expansion versus Unexpanded Indexing We evaluated post-translation document expansion in the 1999 Topic Detection and Tracking evaluation (TDT) in the topic tracking task. In this task, the system is presented with one or more exemplar stories – a form of query-by-example – and must determine whether each new story in an incoming stream is on the same *topic* as the exemplar(s) ³ The TDT-1999 tracking task is multi-lingual, searching stories in both English and Mandarin Chinese, and multi-modal, involving both newswire text and broadcast news audio. We focus on the cross-language component of the tracking task, using English exemplars to identify on-topic stories in Mandarin Chinese. We compare top-1 translation of the Mandarin Chinese stories with and without post-translation document expansion⁴. We use the earlier TDT-2 English newswire text collection as our side collection for expansion. We perform topic tracking on 60 topics with 4 exemplars each. Here, we report the mean average precision on the 55 topics for which there are on-topic Mandarin audio stories. The mean average precision for retrieval of unexpanded documents is 0.3595 while post-translation document expansion raises this figure to 0.41. This difference is significant at $p < 0.01$ by paired t-test, two-tailed. The contrast is illustrated in Figure 3 below.

³A topic is defined as a seminal event or activity along with all directly related events and activities.

⁴Since Mandarin Chinese has little surface morphology, we omit the stemming backoff phase.




Figure 3: Comparison of effectiveness of post-translation document expansion versus bare top-1 translation of Chinese documents: Bars above x-axis indicate document expansion outperforms bare translation

4.2 Overall Retargeting Results

We have found that retargeting can be accomplished quite quickly (a day without document expansion, three days for TREC-sized collections with document expansion), but only if the required infrastructure is in place. Adapting a system that was developed initially for Chinese to handle French documents required several weeks, with most of that effort invested in development of four-stage backoff translation and statistical stemming. Further adapting the system to handle German documents revealed the importance of compound splitting, a problem that we will ultimately need to address by incorporating a more general segmentation strategy than we used initially for Chinese. In extending the system to Italian we have found that although our statistical stemmer presently performs poorly in that language, we can achieve quite credible results even with a fairly small (17,313 term) bilingual term list using the freely available Muscat stemmer (which exist for ten languages). So although it is possible in concept to retarget to a new language in just a few days, extending the system typically takes us between one and three weeks because we are still climbing the learning curve.

5 Conclusion

By building on the lessons learned using the TREC, CLEF, NTCIR, and TDT collections, we have sought to build an infrastructure that can be applied to a broad array of languages. Arabic and Korean collections are expected to become available in the next year, and we are now evolving our interface to support user studies. Our approach is distinguished by support for interactive retrieval even in languages for which machine translation is presently unavailable, and our ultimate goal is to characterize how closely we can approximate the retrieval effectiveness users would obtain if they had the best available machine translations for the retrieved documents.

References

- [1] Douglas W. Oard and Philip Resnik. Support for interactive document selection in cross-language information retrieval. *Information Processing and Management*, 35(3):363–379, July 1999.
- [2] Amit Singhal, John Choi, Donald Hindle, Julia Hirschberg, Fernando Pereira, and Steve Whittaker. AT&T at TREC-7 SDR Track. In *Proceedings of the DARPA Broadcast News Workshop*, 1999.