

Recovering translation errors in cross-language image retrieval using word association models

Masashi Inoue
National Institute of Informatics
m-inoue@nii.ac.jp

Abstract

Connecting short queries and short titles of relevant images is difficult. Different lexical expressions may be used in queries and captions that refer to the same concept. In the ImageCLEF2005 ad hoc task, we investigated the use of learned word association models that represent how pairs of words are related. We compared a precision-oriented simple word-matching retrieval model and a recall-oriented retrieval model with word association models, and we also investigated combinations of models. Experimental results on English and German topics are rather discouraging, as the use of word association models degraded performance. On the other hand, word association models help in retrieval for Japanese topics. Considering the relatively low quality of Japanese-to-English machine translation, this result may indicate that word association could play some role in recovering translation errors at the retrieval stage.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: H.3.3 Information Search and Retrieval

General Terms

Model, Learning, Algorithm, Effectiveness

Keywords

Image retrieval, word association, sparse data, translation error, model combination, result merge

1 Introduction

Retrieval of too many or too few documents for a query causes problems for the users of information retrieval (IR) systems. If given too many documents, they may experience difficulties in finding relevant documents among the results. On the other hand, if given too few documents, there will be little chance of finding relevant information. Of the above two problems, we consider the latter problem: insufficient retrieved results. More precisely, this problem can be divided into two sub-problems: 1) there are not enough relevant documents stored in the database, or 2) relevant documents are not retrieved by the system. We concentrate on the second sub-problem.

Typically in text retrieval, a shortage of retrieved documents is often caused by the problem of term-mismatch: the words in a query do not appear in most documents even though they are relevant to the query. This is essentially the same in ad hoc image retrieval when captions are used as the target of query matching. The difference is that sometimes the number of words in image captions is quite small and term-mismatches are likely to occur more often.

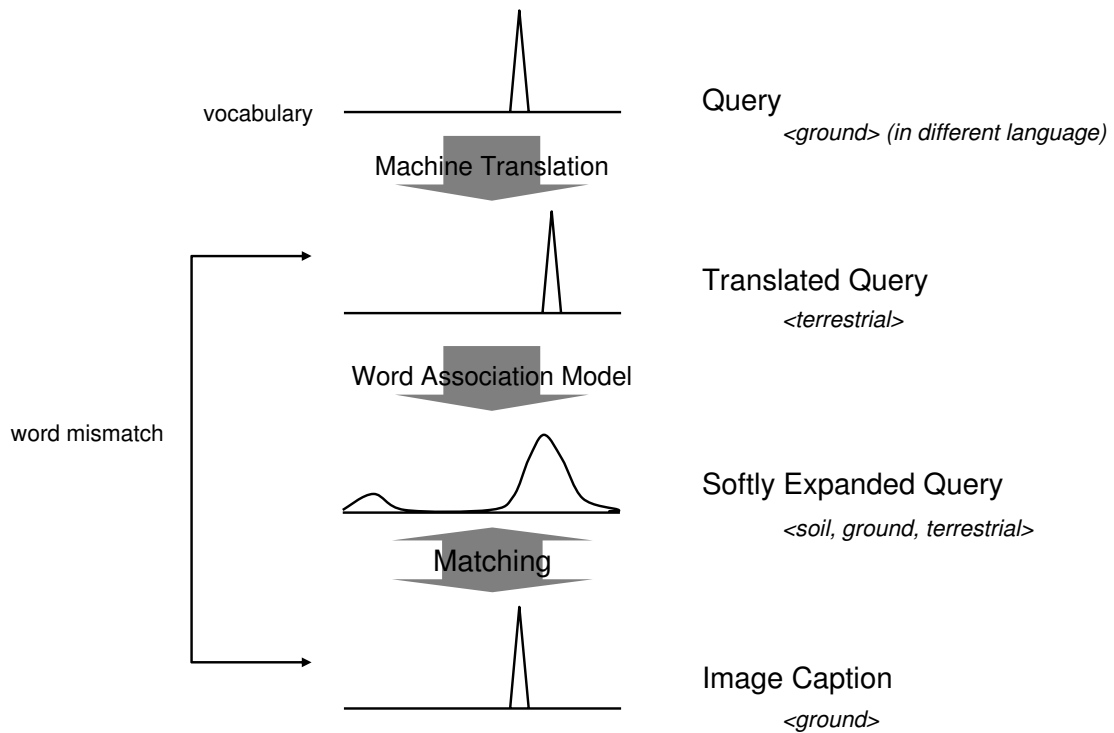


Figure 1: Schematic diagram for the process of recovering translation errors using word association models

One way to mitigate such mismatches is to use an enlarged query word set instead of the small query word set supplied by users. A typical technique is query expansion where some alternative query words are added to the original query from the document set based on relevance or pseudo-relevance judgements.

In ImageCLEF2005, we studied the effects of word association models by employing a kind of probabilistic word-by-word query translation model structure [4], although in our models, the actual translation took place by the MT system outside of the retrieval model. That is, the translation in the model is, in effect, the monolingual word expansion [3]. We tested our approach in the setting where both queries and annotations were short, which are frequently observed characteristics of text-based image retrieval. Concerning the differences between languages, we only considered the influence of machine translation (MT). Monolingual English-to-English, cross-lingual German-to-English, and cross-lingual Japanese-to-English image retrievals were compared. One finding from our experimental runs was that when the simple word matching strategy failed to retrieve relevant images because of erroneous translations, the use of word association models could improve the word matching. The conceptual process of translation error recovery by word association is depicted in Figure 1. In our runs, a recovery effect was observed only in Japanese-to-English translation, an example of translation between disparate languages.

In the following, we first introduce the ImageCLEF2005 image collection and the pre-processing applied to it. Second, we describe the run conditions and retrieval models used. Third, we show the retrieval results on the submitted runs. Finally, we conclude the paper with some discussion.

Table 1: Summary of run conditions

Query Language	English/German/Japanese
Initial Query	title
Query Type	automatic
Feedback/Expansion	with
Modality	text
Document Type	short titles
Translation	machine translation

2 Data Preparation

2.1 Test Collection

The test collection of ImageCLEF2005 consists of 28133 images and their captions in English from the St Andrews Library photographic collection [1]. Each caption has nine fields assigned by experts. These are ‘Record ID’, ‘Short title’, ‘Long title’, ‘Location’, ‘Description’, ‘Date’, ‘Photographer’, ‘Categories’, and ‘Notes’. Such well-annotated images can be found in places such as digital museums and commercial photo collections but are rare in other cases. That is because casual annotators do not have enough knowledge to annotate images systematically, nor they do have any desire to spend time on annotations. For this reason, we are motivated to retrieve less carefully annotated images. Of the fields in the test collection, ‘short titles’ are considered to be the simplest form of annotation. Therefore, we used only short titles for indexing. The mean length of the short titles was 3.43 words. The distribution of lengths had a heavy tail on the short side. The size of the vocabulary was 9883 words for the documents, and 9945 for both documents and queries. The vocabulary contained three irregular words: ‘null’, ‘untitled’, and ‘ ϕ ’, where ϕ is introduced to represent empty titles.

Topics of retrieval were described in three fields: short descriptions (titles), long descriptions (narratives), and example images. In our experiment, we used only short descriptions (titles), which can be regarded as typical queries.

2.2 Run Conditions

The main characteristics of our runs are summarized in Table 1. The most notable point is that we used only the short title field. We were interested more in the exploitation of short text than in utilization of structured and multi-faceted text. Although we were also interested in the use of visual images, we were unable to advance to that level.

Another point we should mention is the use of ‘Feedback/Expansion’. We used only expansion and not feedback. That is, we employed neither manual relevance feedback nor pseudo-relevance feedback. The models for the soft word expansion were built prior to querying and no candidate word selection process was used at the querying stage. The retrieval model we used is explained in 3.1.

The last factor is the query language. We examined English, German, and Japanese. We considered English topics as the baseline, German topics as the relatively ‘easy’ task, and Japanese topics as the relatively ‘hard’ task. Here, by ‘easy’ we mean that the current state-of-the-art accuracy of machine translation is high and retrieval can be conducted in nearly the same fashion as with the source language. Similarly, by ‘hard’, we mean that queries differ substantially from the original ones after going through the machine translation process. According to the results of ImageCLEF2004 that consists of the same image dataset as ImageCLEF2005 but different topics, German topics yielded the highest average MAP score after English, and Japanese topics yielded the lowest average MAP scores for the top five systems [1].

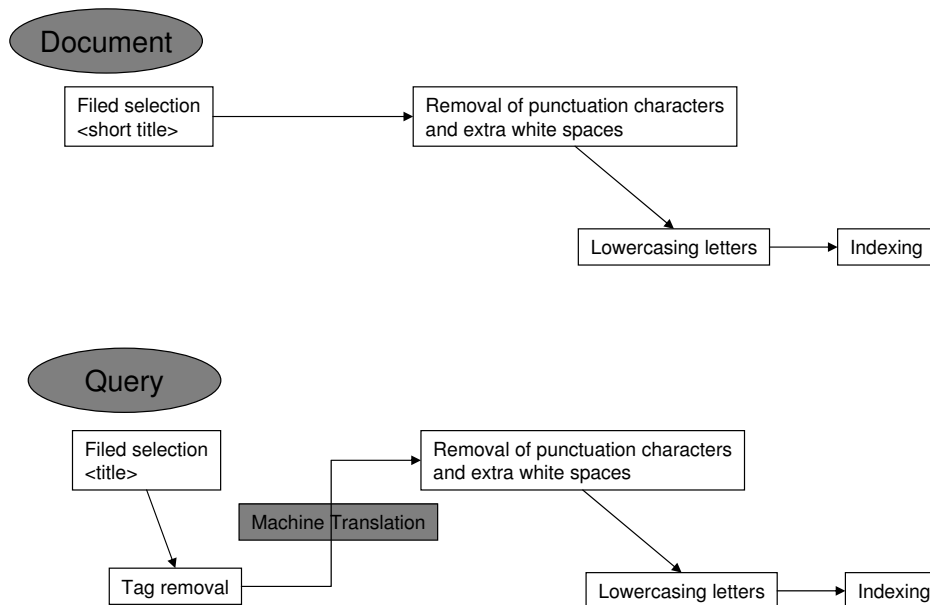


Figure 2: Data pre-processing for documents (image captions) and queries (topics).

2.3 Translation and Pre-Processing

The pre-processing we conducted to prepare data is summarized in Figure 2. As mentioned in the previous section, we used only the title fields of topics and short title fields of captions. Therefore, the initial step was the extraction of those fields from the collection. For topics, titles were surrounded by `<title>` tags; these tags were removed and the bodies of the titles were translated. Although translation is part of the retrieval process, we explain the procedure of translation here because we carried out translations within the process of data preparation. Our approach to cross-language retrieval is query translation. According to previous experiments on ImageCLEF ad hoc data, query translation generally outperforms document translation [2]. We thought that the combination of query translation and document translation might be promising. However, as the starting point, we only consider query translation here. German and Japanese topics were translated into English, the document language, using the Babelfish web-based MT system¹. The translation was done manually: we entered queries in a web form and their respective translations were returned, and punctuation and extra white spaces were removed. The results of translation are shown in Appendix A. Punctuation in the short titles was also removed. Finally, all upper case letters were converted to lower case, and both queries (titles) and documents (short titles) were indexed together.

2.4 Qualitative Analysis of Translation Errors

The results of machine translation usually contain errors that may affect the performance of IR at a subsequent stage, but the relationship is not straightforward. For example, when a word is translated into ‘photographs’ when it should be translated to ‘pictures’, this difference has little effect in understanding sentences that contain the word. Therefore, it may not be considered an error. However, for IR, and image retrieval in particular where only short text descriptions are available, such a difference may change the results of retrieval drastically. For instance, when all relevant images are annotated as ‘pictures’, a query translated as ‘photographs’ cannot retrieve them. On the other hand, when all relevant images are annotated as ‘photographs’, the mistrans-

¹<http://babelfish.altavista.com>

lation benefits the retrieval process. Here, we analyse the results of machine translation of queries from the point of view of their effect on IR.

First, we examine the overall quality of the translations. Translation from German to English was performed well. Among 28 topics (titles), four topics were translated exactly as in the original English – topic numbers 3, 5, 6, and 18 in A.2. This result confirms the relatively high accuracy of German–English MT. Notable errors in German-to-English translation were related to prepositions. For example, ‘at’ in topic 1 should be ‘on’, ‘of’ in topic 12 should be ‘from’, and ‘on’ in topic 14 should be ‘at’. Other typical errors were inappropriate assignment of imprecise synonyms. For example, in topic 1 ‘ground’ is replaced by ‘soil’, and in topics 10 and 28 ‘picture’ is replaced by ‘photographs’. Despite these errors, in most translation results from German, the basic meanings of topics were similar to the original English. More problematic was that three words were not translated into English: ‘Fischer’ in topic 7, ‘puttet’ in topic 15, and ‘Portraitsnahmen’ in topic 26. For simplicity, we treated them as if they were English words. Additionally, the MT system did not translate ‘kutsche’ into any word.

For Japanese-to-English translation, the quality of translation was apparently worse (see A.3). Some of the Japanese words could not be translated at all. Untranslated words were ‘aiona (Iona)’, ‘nabiku (waving)’, and ‘sentoandoryusu (St Andrews)’. The problem here is that the untranslated words were often proper nouns, which might be useful for distinguishing relevant documents from irrelevant documents. Ideally, such out-of-vocabulary words should be translated by using other external sources, such as larger and more up-to-date dictionaries or by transliteration. In this experiment, however, we simply eliminated such untranslated non-ASCII characters from the translation results.

In German-to-English translation, the above two proper names (Iona and St Andrews) could be translated with no problem. This difference can be understood easily by the fact that, in German topics, the words were spelled in the same way as in English. Therefore, no translation was necessary. On the other hand, in Japanese topics, the translator had phonetically converted them from the original English topics to katakana characters. Therefore, untranslated words could not be used as is. Another factor to be considered is that phonetic transcription is not unique. Therefore, even if there was an entry for the word in the dictionary, the back translation by MT systems might not find a relevant entry because of the phonetic ambiguities.

In addition to the above out-of-vocabulary word problems in the MT system, the Japanese-to-English translations contained errors in prepositions similar to those in the German-to-English translations. Errors that were peculiar to the Japanese-to-English translations were the excessive use of definite articles and relative pronouns. We hypothesize that such translations were derived from the design of the MT system, which was designed not for the translation of short phrases such as titles, but for larger units of text such as paragraphs. Thus, the MT system tried to produce *natural* sentences by adding definite articles and relative pronouns to fill the gaps between grammatically disparate languages. Short query translations may require choosing either a sophisticated MT system or simple word-by-word translation, depending on the difficulties of translation.

So far, we have discussed the quality of the topic translations assuming that both German and Japanese topics are equivalent to the English topics in terms of their contents. However, it may be noteworthy that they are the translations from the original English topics. Since relationships between expressions in different languages are not one-to-one, a non-English topic used here was one of many possible translations. Moreover, the expressions in translated queries might not be typical as the queries in that language even if they were correct translations of typical English queries. Therefore, the translation errors analysed above were possibly caused by both machine translation at retrieval stage and by translation ambiguities at topic preparation stage. Although this is not negligible, because it is too involved a subject to be treated here in detail and we can expect the translations by experts were far better than machine translation, we do not consider the influence by the translations when topics were created.

3 Methodology

3.1 Retrieval Models

We introduce retrieval models based on the unigram language model and word association model. The baseline model is the simple keyword matching document model denoted by **diag**. For the query $\mathbf{q} = \{q_1, \dots, q_K\}$, the probability of \mathbf{q} is

$$\prod_{k=1}^K P(q_k | \mathbf{d}_n),$$

where \mathbf{d}_n indicates the n th document or image. For the word association model, we estimated the following transitive probabilities from the j th word to the i th word in the vocabulary:

$$P(w_i | w_j).$$

When the above two models are combined, the following model represents the process of query generation:

$$\prod_{k=1}^K \sum_{i=1}^V P(q_k | w_i) P(w_i | \mathbf{d}_n).$$

Here, we assume independence between query words: $P(\mathbf{q}) = \prod_{k=1}^K P(q_k)$, although this is not always true for the ImageCLEF2005 topics, where titles are sometimes sentential and word orders have meaning.

The word association models can be estimated in various ways, disregarding the statistical justification. We tried three methods. In all three methods, we regarded the frequencies of co-occurrence of two words as the measure of word association. If two words co-occurred, they were assumed to be related. The first method counts self co-occurrences, where a word is regarded as co-occurring with itself, as well as co-occurrences. Values for each term pair are estimated as follows:

$$P(w_i | w_j) = \frac{\#(w_i, w_j)}{\#(w_j)} \quad \text{where } i \neq j \text{ and } \#(w_j) > 0,$$
$$P(w_i | w_i) = \frac{\#(w_i, w_i) + \#(w_i)}{\#(w_i)} \quad \text{where } \#(w_i) > 0.$$

Here, $\#(w_i, w_j)$ represents the frequency of co-occurrence of w_i and w_j (appearance of the two words in the same image caption), and $\#(w_j)$ represents the frequency of occurrence of w_j . This procedure strengthens self-similarities in the model and is termed **cooc**. The second method counts purely co-occurring pairs and is named **coocp**. Values for each term pair are estimated as follows:

$$P(w_i | w_j) = \frac{\#(w_i, w_j)}{\#(w_j)}.$$

The third method normalizes the frequencies of co-occurrences (w_i and w_j) by the frequencies of the word w_j :

$$P(w_i | w_j) = \frac{\#(w_i, w_j)}{\#(w_i)\#(w_j)}.$$

This method is termed **cooct**. The baseline model that does not use word association models can be interpreted as using a diagonal word association model with non-zero elements that are all one. This is why we denoted it as **diag**.

Note that these models were estimated prior to the arrival of queries and the computation at query time focused on score calculation.

3.2 Scoring Functions

Our runs were divided into two groups according to the scoring function employed. In the first group, documents were ranked according to the query–log likelihood of document models. As we used unigram language models for each document, the scoring function for the n th document given the query \mathbf{q} is written as:

$$\log L = \sum_{k=1}^K \log P(q_k | \mathbf{d}_n),$$

where K is the length of the query. When a word association model is used, the function becomes

$$\log L = \sum_{k=1}^K \log \sum_{i=1}^V P(q_k | w_i) P(w_i | \mathbf{d}_n),$$

where V is the vocabulary size. Runs based on these functions were marked with `log_lik`.

In the second group of runs, documents were ranked according to the accumulated information for all matched words. First, we transform the variable for the probability of query word q_k , $P(q)$, to $F_q = e^{(\log P(q))^{-1}}$ where $P(q)$ is either $P(q | \mathbf{d}_n)$ or $\sum_{i=1}^V P(q | w_i) P(w_i | \mathbf{d}_n)$ and is considered only when $P(q) \neq 0$. Then, the new scoring function can be defined as:

$$\log L' = \sum_{k=1}^K \log \frac{1}{F_{q_k}}.$$

We regard $\log \frac{1}{F_{q_k}}$ as the information on query word q . A document with a higher score is assumed to have more information on the query. In general, when an expansion method is involved, the number of terms matched between queries and documents increases. Consequently, the scores of documents given by the first scoring measure `log_lik` are larger in models with expansion than in those without. Thus, the first scoring measure is not suited for the comparison of output scores between different models. The second measure was derived heuristically and is intended to allow combining the outputs of different models. Runs based on this measure were marked with `vt_info`.

3.3 Model Output Combination

When the `vt_info` measure is used, the combination of different models at the output level is simple because their scores are directly comparable. First, two sets of document scores and corresponding document indices from two models are merged. Then, they are sorted in descending order of scores. For each document, the higher score is retained. This assumes that lower scores usually correspond to lack of knowledge about the documents and are thus less reliable. From the eventual rank, the top M documents will be extracted as the final result.

4 Experimental Results

We submitted 16 runs (files), consisting of eight for English, four for German, and four for Japanese topics. We used `wam` as the group name. The names of submission were formed as the concatenation of `group name`, `scoring function`, `word association model`, and `query language`. Regarding the word association model, `dc` represents the combination of `diag` and `cooc` using the method described in 3.3. For English topics, the following runs were submitted:

```
wam_log_lik_diag_e
wam_log_lik_cooc_e
wam_log_lik_coocp_e
```

Table 2: Summary of mean average precision scores for the submitted runs

Method	Scoring Function					
	log-lik			vt-info		
	English	German	Japanese	English	German	Japanese
diag	0.0263	0.0160	0.0062	0.0108	0.0082	0.0057
cooc	0.0161	0.0045	0.0085	0.0089	0.0032	0.0088
coocp	0.0049	N.A.	N.A.	0.0010	N.A.	N.A.
cooct	N.A.	N.A.	N.A.	0.0089	N.A.	N.A.
dc	N.A.	N.A.	N.A.	0.0107	N.A.	N.A.
Runs by other participants						
Highest submitted	0.4135	0.3375	0.2811			
Similar setting	0.1798					

wam_vt_info_diag_e
wam_vt_info_cooc_e
wam_vt_info_coocp_e
wam_vt_info_cooct_e
wam_vt_info_dc_e

For German topics:

wam_log_lik_diag_g
wam_log_lik_cooc_g
wam_vt_info_diag_g
wam_vt_info_cooc_g

For Japanese topics:

wam_log_lik_diag_j
wam_log_lik_cooc_j
wam_vt_info_diag_j
wam_vt_info_cooc_j

Each file contained 980 scores (the 35 top scores for each of 28 topics). We made a mistake when creating the submission files and could not obtain any meaningful official results for these submissions. The figures of mean average precision (MAP) scores in Table 2 are based on the runs we intended to submit. They were calculated after we received the list of relevant images for each topic (qrel file). Overall, our retrieval performances were insufficient. For comparison, we included the MAP scores of the best runs from other participants, as shown at the bottom of Table 2. They are CUHK-ad-eng-tv-kl-jm2 for English topics, R2D2vot2Ge for German topics, and AlCimg05Exp3Jp for Japanese topics. For English runs, we also cited the MAP score of an example run imirt0baset0enen with run conditions similar to ours (title query and short title index).

Having observed the difference between our runs and others, we may now turn to the analysis of our own runs. In English, our best run was actually the `diag` model, which we had considered as the simplest baseline. In contrast, all models with word association underperformed. There are two possible explanations for this result. First, there was no need to relax the limitation of exact term matching. Some relevant documents could be retrieved by word-by-word correspondence and other relevant documents could never be reached by word-level expansion. Second, the word association models were not learned adequately, so they could not help with connecting query words and document words. To clarify which of these two reasons led to this result, we must analyse the data set further. This relationship between the `diag` model and other `cooc`-type models was the same for the German topics. When the `vt_info` scoring function was used, an

important observation is that the MAP scores for `cooc` and `cooct` were the same and those for `diag` and `dc` were nearly the same. By analysing the performances for individual topics, we found that `cooc` and `cooct` behaved in the same way.

Further analysis is required to understand this phenomenon. For `dc`, by analysing the influences of the two models, we observed that the `diag` model dominated the top scores. We had expected this tendency, because an exact-matching scheme should have higher confidence in its outputs when queries can find their counterparts. What was unexpected was that the dominance of the `diag` model often ranged from the top rank to about the 1000th rank, and scores given by `cooc` models appeared only in lower ranks. Even though the ranking was determined by the interlaced ranks from both models, because we had submitted only the top 35 ranks, the resulting MAP scores were determined almost solely by the `diag` model. This outcome was not desirable. For topics 2 and 18, the `cooc` models worked better than the `diag` models. Nevertheless, as explained above, the benefit of the `cooc` models was not taken into account in the final results of the `dc` method. We must consider a better way of rank merging so as not to miss such opportunities.

As we can see in Table 2, the trends of model discrepancy were similar in English and German topics. However, in Japanese topics, the use of word association models (`cooc`) improved performance in both scoring functions. For an explanation of this reversal effect, we can consider the quality of translation. In the `diag` model, when English and Japanese topics were compared, the retrieval performance simply degraded as the translation quality degraded. In contrast, the word association models might provide some improvement. It may be considered as recovering from the translation errors that caused mismatches in the retrieval process. However, the relationship between translation quality and the positive effects of word association models was not simple because it was not monotonic. When comparing `diag` and `cooc` in English and German topics, even though German topics contained some translation errors, the degradation of performance by using `cooc` was severer in German than in English. This problem may be better understood by considering additional languages with varying translation difficulties.

5 Discussion

In our runs, we observed that the use of word association models might help recover query translation errors given by MT systems. However, because our system performed quite poorly in terms of MAP scores, it is difficult to generalize this finding. We need to improve the baseline models to some reasonable level. In the pre-processing and the retrieval models we employed, we did not consider the following three factors that are important to IR performance: 1) idf, 2) stop words, and 3) document length. Incorporation of these factors into the modelling may be the first step towards obtaining a reasonable performance.

If the use of word association models in cross-language retrieval is beneficial for mitigating the effect of translation errors, a similar effect should be observed in other types of expansion techniques. Although we do not know the details of expansion techniques used by other participants, it seems that the use of ‘Expansion/Feedback’ techniques improved performance in most languages. We would like to see if these expansion techniques at query time serve as a more powerful component of retrieval when translation is erroneous than when translation is error free.

Another direction of interest lies in the design of MT systems. In our runs, we used an MT system with a single output. If we had used an MT system with multiple candidate outputs with their confidence scores, the system would have performed the soft expansion by itself. It is not clear whether using such MT systems with our models will improve or degrade the retrieval results.

6 Conclusions and Future Work

Text-based image retrieval that relies on short descriptions such as titles is considered to be less robust to translation errors. In the experiments on the ad hoc task in ImageCLEF2005, word association models helped with the retrieval of Japanese topics when translation into English

using the machine translation system was quite erroneous. We hypothesize that this could be explained by the recovery effect given by word expansion. The above argument might be verified by comparing various languages with different degree of difficulties in English translation.

Two important extensions we could not investigate were the utilization of visual information and the exploitation of training data sets. We are particularly interested in how the use of these will help retrieval for difficult topics in which visual or contextual information plays a vital role.

References

- [1] P. Clough, H. Mueller, and M. Sanderson. The CLEF cross language image retrieval track (ImageCLEF) 2004. In *ImageCLEF2004 Working Note*, 2004.
- [2] Paul Clough. Caption vs. query translation for cross-language image retrieval. In *ImageCLEF2004 Working Note*, 2004.
- [3] Masashi Inoue and Naonori Ueda. Retrieving lightly annotated images using image similarities. In *SAC '05: Proceedings of the 2005 ACM symposium on Applied computing*, pages 1031–1037, NY, USA, March 2005.
- [4] Wessel Kraaij and Franciska de Jong. Transitive CLIR models. In *RIAO*, pages 69–81, Vaucluse, France, April 26-28 2004.

A Results of query translation

A.1 English (Original) Queries

- 1 Aircraft on the ground
- 2 People gathered at bandstand
- 3 Dog in sitting position
- 4 Steam ship docked
- 5 Animal statue
- 6 Small sailing boat
- 7 Fishermen in boat
- 8 Building covered in snow
- 9 Horse pulling cart or carriage
- 10 Sun pictures Scotland
- 11 Swiss mountain scenery
- 12 Postcards from Iona Scotland
- 13 Stone viaduct with several arches
- 14 People at the marketplace
- 15 Golfer putting on green
- 16 Waves breaking on beach
- 17 Man or woman reading
- 18 Woman in white dress
- 19 Composite postcards of Northern Ireland
- 20 Royal visit to Scotland not Fife
- 21 Monument to poet Robert Burns
- 22 Building with waving flag
- 23 Tomb inside church or cathedral
- 24 Close-up picture of bird
- 25 Arched gateway
- 26 Portrait pictures of mixed sex group
- 27 Woman or girl carrying basket
- 28 Colour pictures of woodland scenes around St Andrews

A.2 German to English Translations

- 1 Airplane at the soil
- 2 People meets at the pavilion
- 3 Dog in sitting position
- 4 Docked steam ship
- 5 Animal statue
- 6 Small sailing boat
- 7 Fischer in the boat
- 8 Snow-covered building
- 9 Horse pulls kutsche or car
- 10 Sun photographs Scotland
- 11 Swiss mountain country shank
- 12 Postcards of Iona Scotland
- 13 Stone viaduct with several elbows
- 14 People on the market
- 15 Golf player puttet on the green
- 16 At the beach breaking waves
- 17 Reading man or reading woman
- 18 Woman in white dress
- 19 Postcards with several opinions from Northern Ireland
- 20 Royal attendance in Scotland except Fife
- 21 Monument of the poet Robert Burns
- 22 Building with blowing flag
- 23 Tomb in church or cathedral
- 24 Close-up of a bird
- 25 Arc-shaped gate
- 26 Portraitaufnahmen of mix-sexual groups
- 27 Woman or girl carries basket
- 28 Color photographs of forest landscapes around pc Andrew

A.3 Japanese to English Translations

- 1 Terrestrial airplane
- 2 The people who meet in the field music hall
- 3 The dog which sits down
- 4 The steam ship which is docked to the pier
- 5 Image of animal
- 6 Small-sized sailing ship
- 7 Fishermen on boat
- 8 The building which the snow accumulated
- 9 The horse which pulls the load carriage and the carriage
- 10 Photograph of sun Scotland
- 11 The Swiss mountain scenery
- 12 The illustrated postcards of Scotland and island
- 13 The elevated bridge of the stonework which is plural arch
- 14 People of market
- 15 The golfer who does the pad with the green
- 16 The wave which washes in the beach
- 17 The man or the woman who reads
- 18 Woman of white dress
- 19 Illustrated postcards of the synthesis of province
- 20 The Scottish visit of king family other than fife
- 21 Poet Robert Burns's monument
- 22 Flag building
- 23 Grave inside church and large saintly hall
- 24 Close-up photograph of bird
- 25 Gate of arch type
- 26 Portrait photograph of man and woman mixed group
- 27 The woman or the girl who has the basket
- 28 Color picture of forest scenery of every place