Low-cost, High-performance Translation Retrieval: Dumber is Better

Timothy Baldwin

Department of Computer Science

Tokyo Institute of Technology

2-12-1 O-okayama, Meguro-ku, Tokyo 152-8552 JAPAN

tim@cl.cs.titech.ac.jp

Abstract

In this paper, we compare the relative effects of segment order, segmentation and segment contiguity on the retrieval performance of a translation memory system. We take a selection of both bag-of-words and segment order-sensitive string comparison methods, and run each over both character- and word-segmented data, in combination with a range of local segment contiguity models (in the form of N-grams). Over two distinct datasets, we find that indexing according to simple character bigrams produces a retrieval accuracy superior to any of the tested word N-gram models. Further, in their optimum configuration, bag-of-words methods are shown to be equivalent to segment order-sensitive methods in terms of retrieval accuracy, but much faster. We also provide evidence that our findings are scalable.

1 Introduction

Translation memories (TMs) are a list of **translation records** (source language strings paired with a unique target language translation), which the TM system accesses in suggesting a list of target language (L2) **translation candidates** for a given source language (L1) input (Trujillo, 1999; Planas, 1998). **Translation retrieval** (TR) is a description of this process of selecting from the TM a set of translation records (TRecs) of maximum L1 similarity to a given input. Typically in example-based machine translation, either a single TRec is retrieved from the TM based on a match with the overall L1 input, or the input is partitioned into coherent segments, and individual translations retrieved for each (Sato and Nagao, 1990; Nirenburg et al., 1993); this is the first step toward generating a customised translation for the input. With stand-alone TM systems, on the other hand, the system selects an arbitrary number of translation candidates falling within a certain empirical corridor of similarity with the overall input string, and simply outputs these for manual manipulation by the user in fashioning the final translation.

A key assumption surrounding the bulk of past TR research has been that the greater the match stringency/linguistic awareness of the retrieval mechanism, the greater the final retrieval accuracy will become. Naturally, any appreciation in

retrieval complexity comes at a price in terms of computational overhead. We thus follow the lead of Baldwin and Tanaka (2000) in asking the question: what is the empirical effect on retrieval performance of different match approaches? Here, retrieval performance is defined as the combination of retrieval speed and accuracy, with the ideal method offering fast response times at high accuracy.

In this paper, we choose to focus on retrieval performance within a Japanese-English TR context. One key area of interest with Japanese is the effect that *segmentation* has on retrieval performance. As Japanese is a non-segmenting language (does not explicitly delimit words orthographically), we can take the brute-force approach in treating each string as a sequence of characters (**character-based indexing**), or alternatively call upon segmentation technology in partitioning each string into words (**word-based indexing**). Orthogonal to this is the question of sensitivity to *segment order*. That is, should our match mechanism treat each string as an unorganised multiset of terms (the **bag-of-words** approach), or attempt to find the match that best preserves the original segment order in the input (the **segment order-sensitive** approach)? We tackle this issue by implementing a sample of representative bag-of-words and segment order-sensitive methods and testing the retrieval performance of each. As a third orthogonal parameter, we consider the effects of *segment contiguity*. That is, do matches over contiguous segments provide closer overall translation correspondence than matches over displaced segments? Segment contiguity is either explicitly modelled within the string match mechanism, or provided as an add-in in the form of segment N-grams.

To preempt the major findings of this paper, over a series of experiments we find that character-based indexing is consistently superior to word-based indexing. Furthermore, the bag-of-words methods we test are equivalent in retrieval accuracy to the more expensive segment order-sensitive methods, but superior in retrieval speed. Finally, segment contiguity models provide benefits in terms of both retrieval accuracy and retrieval speed, particularly when coupled with character-based indexing. We thus provide clear evidence that high-performance TR is achievable with naive methods, and moreso that such methods outperform more intricate, expensive methods. That is, the *dumber* the retrieval mechanism, the *better*.

Below, we review the orthogonal parameters of segmentation, segment order and segment contiguity (§ 2). We then present a range of both bag-

of-words and segment order-sensitive string comparison methods (§ 3) and detail the evaluation methodology (§ 4). Finally, we evaluate the different methods in a Japanese–English TR context (§ 5), before concluding the paper (§ 6).

2 Basic Parameters

In this section, we review three parameter types that we suggest impinge on TR performance, namely segmentation, segment order, and segment contiguity.

2.1 Segmentation

Despite non-segmenting languages such as Japanese not making use of segment delimiters, it is possible to artificially partition off a given string into constituent morphemes through the process of **segmentation**. We will collectively term the resultant segments as **words** for the remainder of this paper.

Looking to past research on string comparison methods for TM systems, almost all systems involving Japanese as the source language rely on segmentation (Nakamura, 1989; Sumita and Tsutsumi, 1991; Kitamura and Yamamoto, 1996; Tanaka, 1997), with Sato (1992) and Sato and Kawase (1994) providing rare instances of character-based systems. This is despite Fujii and Croft (1993) providing evidence from Japanese information retrieval that character-based indexing performs comparably to word-based indexing. In analogous research, Baldwin and Tanaka (2000) compared character- and word-based indexing within a Japanese–English TR context and found character-based indexing to hold a slight empirical advantage.

The most obvious advantage of character-based indexing over word-based indexing is that there is no pre-processing overhead. Other arguments for character-based indexing over word-based indexing are that we: (a) avoid the need to commit ourselves to a particular analysis type in the case of ambiguity or unknown words; (b) avoid the need for stemming/lemmatisation; and (c) to a large extent get around problems related to the normalisation of lexical alternation.

Note that all methods described below are applicable to both word- and character-based indexing. To avoid confusion between the two lexeme types, we will collectively refer to the elements of indexing as **segments**.

2.2 Segment Order

Our expectation is that TRecs that preserve the segment order observed in the input string will provide closer-matching translations than TRecs containing those same segments in a different order.

As far as we are aware, there is no TM system operating from Japanese that does not rely on word/segment/character order to some degree. Tanaka (1997) uses pivotal content words identified by the user to search through the TM and locate TRecs which contain those same content words in the same order and preferably the same segment distance apart. Nakamura (1989) similarly gives preference to TRecs in which the content words contained in the original input occur in

the same linear order, although there is the scope to back off to TRecs which do not preserve the original word order. Sumita and Tsutsumi (1991) take the opposite tack in iteratively filtering out NPs and adverbs to leave only functional words and matrix-level predicates, and find TRecs which contain those same key words in the same ordering, preferably with the same segment types between them in the same numbers. Sato and Kawase (1994) employ a more local model of *character* order in modelling similarity according to N-grams fashioned from the original string.

2.3 Segment contiguity

Given the input $\alpha_1\alpha_2\alpha_3\alpha_4$, we would expect that of $\alpha_1\beta_1\alpha_2\beta_2\alpha_3\beta_3\alpha_4$ and $\alpha_1\alpha_2\alpha_3\alpha_4\beta_1\beta_2\beta_3$, the latter would provide a translation more reflective of the translation for the input. This intuition is captured either by embedding some contiguity weighting facility within the string match mechanism (in the case of weighted sequential correspondence — see below), or providing an independent model of segment contiguity in the form of segment N-grams.

The particular N-gram orders we test are simple unigrams (1-grams), pure bigrams (2-grams), and mixed unigrams/bigrams. These N-gram models are implemented as a pre-processing stage, following segmentation (where applicable). All this involves is mutating the original strings into N-grams of the desired order, while preserving the original segment order and segmentation schema. From the Japanese string 夏の雨 [*natu-no-ame*] “summer rain”,¹ for example, we would generate the following variants (common to both character- and word-based indexing):

1-gram:	夏の雨
2-gram:	夏の雨
Mixed 1/2-gram:	夏の夏の雨雨

3 String Comparison Methods

As the starting point for evaluation of the three parameter types targeted in this research, we take two bag-of-words (segment order-oblivious) and three segment order-sensitive methods, thereby modelling the effects of segment order (un)awareness. We then run each method over both segmented and unsegmented data in combination with the various N-gram models proposed above, to capture the full range of parameter settings.

The particular bag-of-word approaches we target are the vector space model (Manning and Schütze, 1999, p300) and “token intersection”. For segment order-sensitive approaches, we test 3-operation edit distance and similarity, and also “weighted sequential correspondence”.

All methods are formulated to operate over an arbitrary *wt* schemata, although in L1 string comparison throughout this paper, we assume that any segment made up entirely of punctuation is given a *wt* of 0, and any other segment a *wt* of 1.

¹ Character boundaries (which double as word boundaries in this case) indicated by “.”.

All methods are subject to a threshold on **translation utility**, and in the case that the threshold is not achieved, the null string is returned. The various thresholds are as follows:

Comparison method	Threshold
Vector space model	0.5
Token intersection	0.4
3-operation edit distance	$len(IN)$
3-operation edit similarity	0.4
Weighted seq. correspondence	0.2

where IN is the input string, and len is the conventional segment length operator.

Various optimisations were made to each string comparison method to reduce retrieval time, of the type described by Baldwin and Tanaka (2000). While the details are beyond the scope of this paper, suffice to say that the segment order-sensitive methods benefited from the greatest optimisation, and that little was done to accelerate the already quick bag-of-words methods.

3.1 Bag-of-Words Methods

Vector Space Model

Within our implementation of the **vector space model** (VSM), the segment content of each string is described as a vector, made up of a single dimension for each segment type occurring within S or T . The value of each vector component is given as the weighted frequency of that type according to its wt value. The string similarity of S and T is then defined as the cosine of the angle between vectors $\vec{S}$ and $\vec{T}$, respectively, calculated as:

$$\cos(\vec{S}, \vec{T}) = \frac{\vec{S} \cdot \vec{T}}{|\vec{S}||\vec{T}|} = \frac{\sum_j s_j t_j}{\sqrt{\sum_j s_j^2} \sqrt{\sum_j t_j^2}}$$

Token Intersection

The **token intersection** of S and T is defined as the cumulative intersecting frequency of segment types appearing in each of the strings, normalised according to the combined segment lengths of S and T using Dice's coefficient. Formally, this equates to:

$$tint(S, T) = \frac{2 \times \sum_{e \in S, T} \min(freq_S(e), freq_T(e))}{len(S) + len(T)}$$

where each e is a segment occurring in either S or T , $freq_S(e)$ is defined as the wt -based frequency of segment type e occurring in string S , and $len(S)$ is the segment length of string S , that is the wt -based count of segments contained in S (similarly for T).

3.2 Segment Order-sensitive Methods

3-op Edit Distance and Similarity

Essentially, the segment-based **3-operation edit distance** between strings S and T is the minimum number of primitive edit operations on single segments required to transform S into T (and

vice versa). The three edit operations are *segment equality* (segments s_i and t_j are identical), *segment deletion* (delete segment s_i) and *segment insertion* (insert segment a into a given position in string S). The cost associated with each operation is determined by the wt values of the operand segments, with the exception of segment equality which is defined to have a fixed cost of 0.

Dynamic programming (DP) techniques are used to determine the minimum edit distance between a given string pair, following the classic 4-operation edit distance formulation of Wagner and Fisher (1974).² For 3-operation edit distance, the edit distance between strings $S = s_1 s_2 \dots s_m$ and $T = t_1 t_2 \dots t_n$ is defined as $D_{3op}(S, T)$:

$$D_{3op}(S, T) = d_3(m, n)$$

$$d_3(i, j) = \begin{cases} 0 & \text{if } i = 0 \wedge j = 0 \\ d_3(0, j-1) + wt(t_j) & \text{if } i = 0 \wedge j \neq 0 \\ d_3(i-1, 0) + wt(s_i) & \text{if } i \neq 0 \wedge j = 0 \\ \min \begin{pmatrix} d_3(i-1, j) + wt(s_i), \\ d_3(i, j-1) + wt(t_j), \\ m_3(i, j) \end{pmatrix} & \text{otherwise} \end{cases}$$

$$m_3(i, j) = \begin{cases} d_3(i-1, j-1) & \text{if } s_i = s_j \\ \infty & \text{otherwise} \end{cases}$$

It is possible to normalise operation edit distance D_{3op} into **3-operation edit similarity** S_{3op} by way of:

$$S_{3op}(S, T) = 1 - \frac{D_{3op}(S, T)}{len(S) + len(T)}$$

Weighted Sequential Correspondence

Weighted sequential correspondence (originally proposed in Baldwin and Tanaka (2000)) goes one step further than edit distance in analysing not only segment sequentiality, but also the contiguity of matching segments.

Weighted sequential correspondence associates an incremental weight (orthogonal to our wt weights) with each matching segment assessing the contiguity of left-neighbouring segments, in the manner described by Sato (1992) for character-based matching. Namely, the k th segment of a matched substring is given the multiplicative weight $\min(k, Max)$, where Max is a positive integer. This weighting up of contiguous matches is facilitated through the DP algorithm given below:

$$S_w(S, T) = s(m, n)$$

$$s(i, j) = \begin{cases} 0 & \text{if } i = 0 \vee j = 0 \\ \max \begin{pmatrix} s(i-1, j), \\ s(i, j-1), \\ s(i-1, j-1) + m_w(i, j) \end{pmatrix} & \text{otherwise} \end{cases}$$

$$m_w(i, j) = \begin{cases} cm(i, j) \times wt(i) & \text{if } s_i = s_j \\ 0 & \text{otherwise} \end{cases}$$

$$cm(i, j) = \begin{cases} 0 & \text{if } i = 0 \vee j = 0 \vee s_i \neq t_j \\ \min(Max, cm(i-1, j-1) + 1) & \text{otherwise} \end{cases}$$

² The fourth operator in 4-operation edit distance is segment substitution.

The final similarity is determined as:

$$WSC(S, T) = \frac{2 \times S_w(S, T)}{len_{WSC}(S) + len_{WSC}(T)}$$

where $len_{WSC}(S)$ is the weighted length of S , defined as:

$$len_{WSC}(S) = \sum_{i=1}^m wt(s_i) \times \min(Max, i)$$

4 Evaluation Specifications

4.1 Details of the Dataset

As our main dataset, we used 3033 unique Japanese–English TRecs extracted from construction machinery field reports for the purposes of this research. Most TRecs comprise a single sentence, with an average Japanese character length of 27.7 and English word length of 13.3. Importantly, our dataset constitutes a controlled language, that is, a given word will tend to be translated identically across all usages, and only a limited range of syntactic constructions are employed.

In secondary evaluation of retrieval performance over differing data sizes, we extracted 61,236 Japanese–English TRecs from the JEIDA parallel corpus (Isahara, 1998), which is made up of government white papers. The alignment granularity of this second corpus is much coarser than for the first corpus, with a single TRec often extending over multiple sentences. The average Japanese character length of each TRec is 76.3, and the average English word length is 35.7. The language used in the JEIDA corpus is highly constrained, although not as controlled as that in the first corpus.

The construction of TRecs from both corpora was based on existing alignment data, and no further effort was made to subdivide partitions.

For Japanese word-based indexing, segmentation was carried out primarily with ChaSen v2.0 (Matsumoto et al., 1999), and where specifically mentioned, JUMAN v3.5 (Kurohashi and Nagao, 1998) and ALTJAWS³ were also used.

4.2 Semi-stratified Cross Validation

Retrieval accuracy was determined by way of 10-fold semi-stratified cross validation over the dataset. As part of this, all Japanese strings of length 5 characters or less were extracted from the dataset, and cross validation was performed over the residue, including the shorter strings in the training data (i.e. TM) on each iteration.

In N-fold stratified cross validation, the dataset is divided into N equally-sized partitions of uniform class distribution. Evaluation is then carried out N times, taking each partition as the held-out test data, and the remaining partitions as the training data on each iteration; the overall accuracy is averaged over the N data configurations. As our dataset is not pre-classified according to a discrete class description, we are not able to perform true data stratification over the class distribution. Instead, we carry out “semi-stratification” over the L1 segment lengths of the TRecs.

³ <http://www.kecl.ntt.co.jp/icl/mtg/resources/altjaws.html>

4.3 Evaluation of the Output

Evaluation of retrieval accuracy is carried out according to a modified version of the method proposed by Baldwin and Tanaka (2000). The first step in this process is to determine the set of “optimal” translations by way of the same basic TR procedure as described above, except that we use the held-out translation for each input to search through the L2 component of the TM. As for L1 TR, a threshold on translation utility is then applied to ascertain whether the optimal translations are similar enough to the model translation to be of use, and in the case that this threshold is not achieved, the empty string is returned as the sole optimal translation.

Next, we proceed to ascertain whether the actual system output coincides with one of the optimal translations, and rate the accuracy of each method according to the proportion of optimal outputs. If multiple outputs are produced, we select from among them randomly. This guarantees a unique translation output and differs from the methodology of Baldwin and Tanaka (2000), who judged the system output to be “correct” if the potentially multiple set of top-ranking outputs contained an optimal translation, placing methods with greater fan-out of outputs at an advantage.

So as to filter out any bias towards a given string comparison method in TR, we determine translation optimality based on both 3-operation edit distance (operating over English word bigrams) and also weighted sequential correspondence (operating over English word unigrams). We then derive the final translation accuracy as the average of the accuracies from the respective evaluation sets. Here again, our approach differs from that of Baldwin and Tanaka (2000), who based determination of translation optimality exclusively on 3-operation edit distance (operating over word unigrams), a method which we found to produce a strong bias toward 3-operation edit distance in L1 TR.

In determining translation optimality, all punctuation and stop words were first filtered out of each L2 (English) string, and all remaining segments scored at a *wt* of 1. Stop words are defined as those contained within the SMART (Salton, 1971) stop word list.⁴

Perhaps the main drawback of our approach to evaluation is that we assume a unique model translation for each input, where in fact, multiple translations of equivalent quality could reasonably be expected to exist. In our case, however, both corpora represent relatively controlled languages and language use is hence highly predictable. The proposed evaluation methodology is thus justified.

5 Results and Supporting Evidence

5.1 Basic evaluation

In this section, we test our five string comparison methods over the construction machinery corpus, under both character- and word-based indexing, and with each of unigrams, bigrams and mixed unigrams/bigrams. The retrieval accuracies and times for the different string comparison methods are presented in Figs. 1 and 2, respectively.

⁴ <http://ftp.cornell.cs.edu/pub/smart/english.stop>

Figure 1: Basic retrieval accuracies

Here and in subsequent graphs, “VSM” refers to the vector space model, “TINT” to token intersection, “3opD” to 3-op edit distance, “3opS” to 3-op edit similarity, and “WSC” to weighted sequential correspondence; the bag-of-words methods are labelled in italics and the segment order-sensitive methods in bold. In Figs. 1 and 2, results for the three N-gram models are presented separately, within each of which, the data is sectioned off into the different string comparison methods. Weighted sequential correspondence was tested with a unigram model only, due to its inbuilt modelling of segment contiguity. Bars marked with an asterisk indicate a statistically significant⁵ gain over the corresponding indexing paradigm (i.e. character-based indexing vs. word-based indexing for a given string comparison method and N-gram order). Times in Fig. 2 are calibrated relative to 3-operation edit distance with word unigrams, and plotted against a logarithmic time axis.

Results to come from these figures can be summarised as follows:

- Character-based indexing is consistently superior to word-based indexing, particularly when combined with bigrams or mixed unigrams/bigrams.
- In terms of raw translation accuracy, there is very little to separate the best of the bag-of-words methods from the best of the segment order-sensitive methods.
- With character-based indexing, bigrams offer tangible gains in translation accuracy at the same time as greatly accelerating the retrieval process. With word-based indexing, mixed unigrams/bigrams offer the best balance of translation accuracy and computational cost.
- Weighted sequential correspondence is moderately successful in terms of accuracy, but grossly expensive.

Based on the above results, we judge bigrams to be the best segment contiguity model for character-based indexing, and mixed unigrams/bigrams to be the best segment contiguity

⁵ As determined by the paired t test ($p < 0.05$).

Figure 2: Basic unit retrieval times

model for word-based indexing, and for the remainder of this paper, present only these two sets of results.

While we have been able to confirm the finding of Baldwin and Tanaka (2000) that character-based indexing is superior to word-based indexing, we are no closer to determining why this should be the case. In the following sections we look to shed some light on this issue by considering each of: (i) the retrieval accuracy for other segmentation systems, (ii) the effects of lexical normalisation, and (iii) the scalability and reproducibility of the given results over different datasets. Finally, we present a brief qualitative explanation for the overall results.

5.2 The effects of segmentation and lexical normalisation

Above, we observed that segmentation consistently brought about a degradation in translation retrieval for the given dataset. Automated segmentation inevitably leads to errors, which could possibly impinge on the accuracy of word-based indexing. Alternatively, the performance drop could simply be caused somehow by our particular choice of segmentation module, that is ChaSen.

First, we used JUMAN to segment the construction machinery corpus, and evaluated the resultant dataset in the exact same manner as for the ChaSen output. Similarly, we ran a development version of ALTJAWS over the same corpus to produce two datasets, the first simply segmented and the second both segmented and lexically normalised. By lexical normalisation, we mean that each word is converted to its canonical form. The main segment types that normalisation has an effect on are verbs and adjectives (conjugating words), and also loan-word nouns with an optional long final vowel (e.g. *monitā* “monitor” $\Rightarrow$ *monita*) and words with multiple inter-replaceable kanji realisations (e.g. 充分 [*zyūbu*] “sufficient” $\Rightarrow$ 十分).

The retrieval accuracies for JUMAN, and ALTJAWS with and without lexical normalisation are presented in Fig. 3, juxtaposed against the retrieval accuracies for character-based indexing (bigrams) and also ChaSen (mixed unigrams/bigrams) from Section 5.1. Asterisked bars

Figure 3: Results using different segmentation modules

indicate a statistically significant gain in accuracy over ChaSen.

Looking first to the results for JUMAN, there is a gain in accuracy over ChaSen for all string comparison methods. With ALTJAWS, also, a consistent gain in performance is evident with simple segmentation, the degree of which is significantly higher than for JUMAN. The addition of lexical normalisation enhances this effect marginally. Notice that character-based indexing (based on character bigrams) holds a clear advantage over the best of the word-based indexing results for all string comparison methods.

Based on the above, we can state that the choice of segmentation system does have a modest impact on retrieval accuracy, but that the effects of lexical normalisation are highly localised. In the following, we look to quantify the relationship between retrieval and segmentation accuracy.

In the next step of evaluation, we took a random sample of 200 TRecs from the original dataset, and ran each of ChaSen, JUMAN and ALTJAWS over the Japanese component of each. We then manually evaluated the output in terms of segment precision and recall, defined respectively as:

$$\text{Segment precision} = \frac{\# \text{ correct segs in output}}{\text{Total } \# \text{ segs in output}}$$

$$\text{Segment recall} = \frac{\# \text{ correct segs in output}}{\text{Total } \# \text{ segs in model data}}$$

One slight complication in evaluating the output of the three systems is that they adopt incongruent models of conjugation. We thus made allowance for variation in the analysis of verb and adjective complexes, and focused on the segmentation of noun complexes.

A performance breakdown for ChaSen (CS), JUMAN (JM) and ALTJAWS (AJ) is presented in Tab. 1. ALTJAWS was found to outperform the remaining two systems in terms of segment precision, while ChaSen and JUMAN performed at the exact same level of segment precision. Looking next to segment recall, ChaSen significantly outperformed both ALTJAWS and JUMAN. The source of almost all errors in recall, and roughly half of errors in precision for both ChaSen and

	CS	JM	AJ
Ave. segs/TRec	13.0	12.0	11.7
Segment precision	98.3%	98.3%	98.6%
Segment recall	98.1%	96.2%	97.7%
Sentence accuracy	70.5%	59.0%	72.0%
Total segment types	650	656	634

Table 1: Segmentation performance

JUMAN was katakana sequences such as *gēto-rokku-barubu* “gate-lock valve”, transcribed from English. ALTJAWS, on the other hand, was remarkably successful at segmenting katakana word sequences, achieving a segment precision of 100% and segment recall approaching 99%. This is thought to have been the main cause for the disparity in retrieval accuracy for the three systems, aggravated by the fact that most katakana sequences were key technical terms.

To gain an insight into consistency in the case of error, we further calculated the total number of segment types in the output, expecting to find a core set of correctly-analysed segments, of relatively constant size across the different systems, plus an unpredictable component of segment errors, of variable size. The system generating the fewest segment types can thus be said to be the most consistent.

Based on the segment type counts in Tab. 1, ALTJAWS errs more consistently than the remaining two systems, and there is very little to separate ChaSen and JUMAN. This is thought to have had some impact on the inflated retrieval accuracy for ALTJAWS.

To summarise, there would seem to be a direct correlation between segmentation accuracy and retrieval performance, with segmentation accuracy on key terms (katakana sequences) having a particularly keen effect on translation retrieval. In this respect, ALTJAWS is superior to both ChaSen and JUMAN for the target domain. Additionally, complementing segmentation with lexical normalisation would seem to produce meager performance gains. Lastly, despite the slight gains to word-based indexing with the different segmentation systems, it is still significantly inferior to character-based indexing.

5.3 Scalability of performance

All results to date have arisen from evaluation over a single dataset of fixed size. In order to validate the basic findings from above and observe how increases in the data size affect retrieval performance, we next ran the string comparison methods over differing-sized subsets of the JEIDA corpus.

We simulate TMs of differing size by randomly splitting the JEIDA corpus into ten partitions, and running the various methods first over partition 1, then over the combined partitions 1 and 2, and so on until all ten partitions are combined together into the full corpus. We tested all string comparison methods other than weighted sequential correspondence over the ten subsets of the JEIDA corpus. Weighted sequential correspondence was excluded from evaluation due to its overall sub-standard retrieval performance. The translation accuracies for the different methods

Figure 4: Retrieval accuracies over datasets of increasing size

over the ten datasets of varying size, are indicated in Fig. 4, with each string comparison method tested under character bigrams (“2-gram –seg”) and mixed word unigrams/bigrams (“1/2-gram +seg”) as above. The results for token intersection have been omitted from the graph due to their being almost identical to those for VSM.

A striking feature of the graph is that it is right-decreasing, which is essentially an artifact of the inflated length of each TRec (see Section 4.1) and resultant data sparseness. That is, for smaller datasets, in the bulk of cases, no TRec in the TM is similar enough to the input to warrant consideration as a translation candidate (i.e. the translation utility threshold is generally not achieved). For larger datasets, on the other hand, we are having to make more subtle choices as to the final translation candidate.

One key trend in Fig. 4 is the superiority of character- over word-based indexing for each of the three string comparison methods, at a relatively constant level as the TM size grows. Also of interest is the finding that there is very little to distinguish bag-of-words from segment order-sensitive methods in terms of retrieval accuracy in their respective best configurations.

As with the original dataset from above, 3-operation edit similarity was the strongest performer just nosing out (character bigram-based) VSM for line honours, with 3-operation edit distance lagging well behind.

Next, we turn to consider the mean unit retrieval times for each method, under the two indexing paradigms. Times are presented in Fig. 5, plotted once again on a logarithmic scale in order to fit the full fan-out of retrieval times onto a single graph. VSM and 3-operation edit distance were the most consistent performers, both maintaining retrieval speeds in line with those for the original dataset at around or under 1.0 (i.e. the same retrieval time per input as 3-operation edit distance run over word unigrams for the construction machinery dataset). Most importantly, only minor increases in retrieval speed were evident as the TM size increased, which were then reversed for the larger datasets. All three string comparison methods displayed this convex shape, although the final running time for 3-operation edit similarity under character- and word-based indexing

Figure 5: Relative unit retrieval times over datasets of increasing size

was, respectively, around 10 and 100 times slower than that for VSM or 3-operation edit distance over the same dataset.

To combine the findings for accuracy and speed, VSM under character-based indexing suggests itself as the pick of the different system configurations, combining both speed and consistent accuracy. That is, it offers the best overall retrieval performance.

5.4 Qualitative evaluation

Above, we established that character-based indexing is superior to word-based indexing for distinct datasets and a range of segmentation modules, even when segmentation is coupled with lexical normalisation. Additionally, we provided evidence to the effect that bag-of-words methods offer superior translation retrieval performance to segment order-sensitive methods. We are still no closer, however, to determining why this should be the case. Here, we seek to provide an explanation for these intriguing results.

First comparing character- and word-based indexing, we found that the disparity in retrieval accuracy was largely related to the scoring of katakana words, which are significantly longer in character length than native Japanese words. For the construction machinery dataset as analysed with ChaSen, for example, the average character length of katakana words is 3.62, as compared to 2.05 overall. Under word-based indexing, all words are treated equally and character length does not enter into calculations. Thus a katakana word is treated identically to any other word type. Under character-based indexing, on the other hand, the longer the word, the more segments it generates, and a single matching katakana sequence thus tends to contribute more heavily to the final score than other words. Effectively, therefore, katakana sequences receive a higher score than kanji and other sequences, producing a preference for TRecs which incorporate the same katakana sequences as the input. As noted above, katakana sequences generally represent key technical terms, and such weighting thus tends to be beneficial to retrieval accuracy.

We next examine the reason for the high correlation in retrieval accuracy between bag-of-words and segment order-sensitive methods in their op-

timum configurations (i.e. when coupled with character bigrams). Essentially, the probability of a given segment set permuting in different string contexts diminishes as the number of co-occurring segments decreases. That is, for a given string pair, the greater the segment overlap between them (relative to the overall string lengths), the lower the probability that those segments are going to occur in different orderings. This is particularly the case when local segment contiguity is modelled within the segment description, as occurs for the character bigram and mixed word uni/bigram models. For high-scoring matches, therefore, segment order sensitivity becomes largely superfluous, and the slight edge in retrieval accuracy for segment order-sensitive methods tends to come for mid-scoring matches, in the vicinity of the translation utility threshold.

6 Conclusion

This research has been concerned with the relative import of segmentation, segment order and segment contiguity on translation retrieval performance. We simulated the effects of word order sensitivity vs. bag-of-words word order insensitivity by implementing a total of five comparison methods: two bag-of-words approaches and three word order-sensitive approaches. Each of these methods was then tested under character-based and word-based indexing and in combination with a range of N-gram models, and the relative performance of each such system configuration evaluated. Character-based indexing was found to be superior to word-based indexing, particularly when supplemented with a character bigram model.

We went on to discover a strong correlation between retrieval accuracy and segmentation accuracy/consistency, and that lexical normalisation produces marginal gains in retrieval performance. We further tested the effects of incremental increases in data on retrieval performance, and confirmed our earlier finding that character-based indexing is superior to word-based indexing. At the same time, we discovered that in their best configurations, the retrieval accuracies of our bag-of-words and segment order sensitive string comparison methods are roughly equivalent, but that the computational overhead for bag-of-words methods to achieve that accuracy is considerably lower than that for segment order sensitive methods.

References

- T. Baldwin and H. Tanaka. 2000. The effects of word order and segmentation on translation retrieval performance. In *Proc. of the 18th International Conference on Computational Linguistics (COLING 2000)*, pages 35–41.
- H. Fujii and W.B. Croft. 1993. A comparison of indexing techniques for Japanese text retrieval. In *Proc. of 16th International ACM-SIGIR Conference on Research and Development in Information Retrieval (SIGIR'93)*, pages 237–46.
- H. Isahara. 1998. JEIDA's English–Japanese bilingual corpus project. In *Proc. of the 1st International Conference on Language Resources and Evaluation (LREC'98)*, pages 471–81.
- E. Kitamura and H. Yamamoto. 1996. Translation retrieval system using alignment data from parallel texts. In *Proc. of the 53rd Annual Meeting of the IPSJ*, volume 2, pages 385–6. (In Japanese).
- S. Kurohashi and M. Nagao. 1998. *Nihongo keitai-kaiseki sisutemu JUMAN* [Japanese morphological analysis system JUMAN] version 3.5. Technical report, Kyoto University. (In Japanese).
- C. Manning and H. Schütze. 1999. *Foundations of Statistical Natural Language Processing*. MIT Press.
- Y. Matsumoto, A. Kitauchi, T. Yamashita, and Y. Hirano. 1999. *Japanese Morphological Analysis System ChaSen Version 2.0 Manual*. Technical Report NAIST-IS-TR99009, NAIST.
- N. Nakamura. 1989. Translation support by retrieving bilingual texts. In *Proc. of the 38th Annual Meeting of the IPSJ*, volume 1, pages 357–8. (In Japanese).
- S. Nirenburg, C. Domashnev, and D.J. Grannes. 1993. Two approaches to matching in example-based machine translation. In *Proc. of the 5th International Conference on Theoretical and Methodological Issues in Machine Translation (TMI-93)*, pages 47–57.
- E. Planas. 1998. *A Case Study on Memory Based Machine Translation Tools*. PhD Fellow Working Paper, United Nations University.
- G. Salton. 1971. *The SMART Retrieval System: Experiments in Automatic Document Processing*. Prentice-Hall.
- S. Sato and T. Kawase. 1994. *A High-Speed Best Match Retrieval Method for Japanese Text*. Technical Report IS-RR-94-9I, JAIST.
- S. Sato and M. Nagao. 1990. Toward memory-based translation. In *Proc. of the 13th International Conference on Computational Linguistics (COLING '90)*, pages 247–52.
- S. Sato. 1992. CTM: An example-based translation aid system. In *Proc. of the 14th International Conference on Computational Linguistics (COLING '92)*, pages 1259–63.
- E. Sumita and Y. Tsutsumi. 1991. A practical method of retrieving similar examples for translation aid. *Transactions of the IEICE*, J74-D-II(10):1437–47. (In Japanese).
- H. Tanaka. 1997. An efficient way of gauging similarity between long Japanese expressions. In *Information Processing Society of Japan SIG Notes*, volume 97, no. 85, pages 69–74. (In Japanese).
- A. Trujillo. 1999. *Translation Engines: Techniques for Machine Translation*. Springer Verlag.
- A. Wagner and M. Fisher. 1974. The string-to-string correction problem. *Journal of the ACM*, 21(1):168–73.