

A Word-to-Word Model of Translational Equivalence

I. Dan Melamed

Dept. of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, 19104, U.S.A.
melamed@unagi.cis.upenn.edu

Abstract

Many multilingual NLP applications need to translate words between different languages, but cannot afford the computational expense of inducing or applying a full translation model. For these applications, we have designed a fast algorithm for estimating a partial translation model, which accounts for translational equivalence only at the word level. The model's precision/recall trade-off can be directly controlled via one threshold parameter. This feature makes the model more suitable for applications that are not fully statistical. The model's hidden parameters can be easily conditioned on information extrinsic to the model, providing an easy way to integrate pre-existing knowledge such as part-of-speech, dictionaries, word order, etc.. Our model can link word tokens in parallel texts as well as other translation models in the literature. Unlike other translation models, it can automatically produce dictionary-sized translation lexicons, and it can do so with over 99% accuracy.

1 Introduction

Over the past decade, researchers at IBM have developed a series of increasingly sophisticated statistical models for machine translation (Brown et al., 1988; Brown et al., 1990; Brown et al., 1993a). However, the IBM models, which attempt to capture a broad range of translation phenomena, are computationally expensive to apply. Table look-up using an explicit translation lexicon is sufficient and preferable for many multilingual NLP applications, including "crummy" MT on the World Wide Web (Church & Hovy, 1993), certain machine-assisted translation tools (e.g. (Macklovitch, 1994; Melamed, 1996b)),

concordancing for bilingual lexicography (Catizone et al., 1993; Gale & Church, 1991), computer-assisted language learning, corpus linguistics (Melby, 1981), and cross-lingual information retrieval (Oard & Dorr, 1996).

In this paper, we present a fast method for inducing accurate translation lexicons. The method assumes that words are translated one-to-one. This assumption reduces the explanatory power of our model in comparison to the IBM models, but, as shown in Section 3.1, it helps us to avoid what we call indirect associations, a major source of errors in other models. Section 3.1 also shows how the one-to-one assumption enables us to use a new greedy competitive linking algorithm for re-estimating the model's parameters, instead of more expensive algorithms that consider a much larger set of word correspondence possibilities. The model uses two hidden parameters to estimate the confidence of its own predictions. The confidence estimates enable direct control of the balance between the model's precision and recall via a simple threshold. The hidden parameters can be conditioned on prior knowledge about the bitext to improve the model's accuracy.

2 Co-occurrence

With the exception of (Fung, 1995b), previous methods for automatically constructing statistical translation models begin by looking at word co-occurrence frequencies in bitexts (Gale & Church, 1991; Kumano & Hirakawa, 1994; Fung, 1995a; Melamed, 1995). A *bitext* comprises a pair of texts in two languages, where each text is a translation of the other. Word co-occurrence can be defined in various ways. The most common way is to divide each half of the bitext into an equal number of segments and to align the segments so that each pair of segments S_i and T_i are translations of each other (Gale & Church, 1991; Melamed, 1996a). Then, two word tokens (u, v) are said to **co-occur** in the

aligned segment pair i if $u \in S_i$ and $v \in T_i$. The co-occurrence relation can also be based on distance in a bitext space, which is a more general representation of bitext correspondence (Dagan et al., 1993; Resnik & Melamed, 1997), or it can be restricted to words pairs that satisfy some matching predicate, which can be extrinsic to the model (Melamed, 1995; Melamed, 1997).

3 The Basic Word-to-Word Model

Our translation model consists of the hidden parameters λ^+ and λ^- , and likelihood ratios $L(\mathbf{u}, \mathbf{v})$. The two hidden parameters are the probabilities of the model generating true and false positives in the data. $L(\mathbf{u}, \mathbf{v})$ represents the likelihood that \mathbf{u} and \mathbf{v} can be mutual translations. For each co-occurring pair of word types \mathbf{u} and \mathbf{v} , these likelihoods are initially set proportional to their co-occurrence frequency $n_{(\mathbf{u}, \mathbf{v})}$ and inversely proportional to their marginal frequencies $n_{(\mathbf{u})}$ and $n_{(\mathbf{v})}$ ¹, following (Dunning, 1993)². When the $L(\mathbf{u}, \mathbf{v})$ are re-estimated, the model's hidden parameters come into play.

After initialization, the model induction algorithm iterates:

1. Find a set of "links" among word tokens in the bitext, using the likelihood ratios and the competitive linking algorithm.
2. Use the links to re-estimate λ^+ , λ^- , and the likelihood ratios.
3. Repeat from Step 1 until the model converges to the desired degree.

The competitive linking algorithm and its one-to-one assumption are detailed in Section 3.1. Section 3.1 explains how to re-estimate the model parameters.

3.1 Competitive Linking Algorithm

The competitive linking algorithm is designed to overcome the problem of indirect associations, illustrated in Figure 1. The sequences of u 's and v 's represent corresponding regions of a bitext. If u_k and v_k co-occur much more often than expected by chance, then any reasonable model will deem them likely to be mutual translations. If u_k and v_k are indeed mutual translations, then their tendency to

¹The co-occurrence frequency of a word type pair is simply the number of times the pair co-occurs in the corpus. However, $n_{(\mathbf{u})} = \sum_{\mathbf{v}} n_{(\mathbf{u}, \mathbf{v})}$, which is *not* the same as the frequency of \mathbf{u} , because each token of \mathbf{u} can co-occur with several different \mathbf{v} 's.

²We could just as easily use other symmetric "association" measures, such as ϕ^2 (Gale & Church, 1991) or the Dice coefficient (Smadja, 1992).




Figure 1: u_k and v_k often co-occur, as do u_k and u_{k+1} . The direct association between u_k and v_k , and the direct association between u_k and u_{k+1} give rise to an indirect association between v_k and u_{k+1} .

co-occur is called a **direct association**. Now, suppose that u_k and u_{k+1} often co-occur within their language. Then v_k and u_{k+1} will also co-occur more often than expected by chance. The arrow connecting v_k and u_{k+1} in Figure 1 represents an **indirect association**, since the association between v_k and u_{k+1} arises only by virtue of the association between each of them and u_k . Models of translational equivalence that are ignorant of indirect associations have "a tendency ... to be confused by collocates" (Dagan et al., 1993).

Fortunately, indirect associations are usually not difficult to identify, because they tend to be weaker than the direct associations on which they are based (Melamed, 1996c). The majority of indirect associations can be filtered out by a simple competition heuristic: Whenever several word tokens u_i in one half of the bitext co-occur with a particular word token v in the other half of the bitext, the word that is most likely to be v 's translation is the one for which the likelihood $L(\mathbf{u}, \mathbf{v})$ of translational equivalence is highest. The competitive linking algorithm implements this heuristic:

1. Discard all likelihood scores for word types deemed unlikely to be mutual translations, i.e. all $L(\mathbf{u}, \mathbf{v}) < 1$. This step significantly reduces the computational burden of the algorithm. It is analogous to the step in other translation model induction algorithms that sets all probabilities below a certain threshold to negligible values (Brown et al., 1990; Dagan et al., 1993; Chen, 1996). To retain word type pairs that are at least twice as likely to be mutual translations than not, the threshold can be raised to 2. Conversely, the threshold can be lowered to buy more coverage at the cost of a larger model that will converge more slowly.
2. Sort all remaining likelihood estimates $L(\mathbf{u}, \mathbf{v})$ from highest to lowest.
3. Find \mathbf{u} and \mathbf{v} such that the likelihood ratio $L(\mathbf{u}, \mathbf{v})$ is highest. Token pairs of these types

$n_{(u,v)}$	=	frequency of co-occurrence between word types u and v
N	=	$\sum_{(u,v)} n_{(u,v)}$ = total number of co-occurrences in the bitext
$k_{(u,v)}$	=	frequency of links between word types u and v
K	=	$\sum_{(u,v)} k_{(u,v)}$ = total number of links in the bitext
τ	=	$\Pr(\text{mutual translations} \mid \text{co-occurrence})$
λ	=	$\Pr(\text{link} \mid \text{co-occurrence})$
λ^+	=	$\Pr(\text{link} \mid \text{co-occurrence of mutual translations})$
λ^-	=	$\Pr(\text{link} \mid \text{co-occurrence of not mutual translations})$
$B(k n, p)$	=	$\Pr(k n, p)$, where k has a binomial distribution with parameters n and p
N.B.: λ^+ and λ^- need not sum to 1, because they are conditioned on different events.		

Figure 2: Variables used to estimate the model parameters.

would be the winners in any competitions involving u or v .

4. Link all token pairs (u, v) in the bitext.
5. The one-to-one assumption means that linked words cannot be linked again. Therefore, remove all linked word tokens from their respective texts.
6. If there is another co-occurring word token pair (u, v) such that $L(u, v)$ exists, then repeat from Step 3.

The competitive linking algorithm is more greedy than algorithms that try to find a *set* of link types that are jointly most probable over some segment of the bitext. In practice, our linking algorithm can be implemented so that its worst-case running time is $O(lm)$, where l and m are the lengths of the aligned segments.

The simplicity of the competitive linking algorithm depends on the **one-to-one assumption**: Each word translates to at most one other word. Certainly, there are cases where this assumption is false. We prefer not to model those cases, in order to achieve higher accuracy with less effort on the cases where the assumption is true.

3.2 Parameter Estimation

The purpose of the competitive linking algorithm is to help us re-estimate the model parameters. The variables that we use in our estimation are summarized in Figure 2. The linking algorithm produces a set of links between word tokens in the bitext. We define a **link token** to be an ordered pair of word tokens, one from each half of the bitext. A **link type** is an ordered pair of word types. Let $n_{(u,v)}$ be the co-occurrence frequency of u and v and $k_{(u,v)}$ be the number of links between tokens of u and v ³. An

³Note that $k_{(u,v)}$ depends on the linking algorithm, but $n_{(u,v)}$ is a constant property of the bitext.

important property of the competitive linking algorithm is that the ratio $k_{(u,v)}/n_{(u,v)}$ tends to be very high if u and v are mutual translations, and quite low if they are not. The bimodality of this ratio for several values of $n_{(u,v)}$ is illustrated in Figure 3. This figure was plotted after the model's first iteration over 300000 aligned sentence pairs from the




Figure 3: A fragment of the joint frequency $(k_{(u,v)}/n_{(u,v)}, n_{(u,v)})$. Note that the frequencies are plotted on a log scale — the bimodality is quite sharp.

Canadian Hansard bitext. Note that the frequencies are plotted on a log scale — the bimodality is quite sharp.

The linking algorithm creates all the links of a given type independently of each other, so the number $k_{(u,v)}$ of links connecting word types u and v has a binomial distribution with parameters $n_{(u,v)}$ and $p_{(u,v)}$. If u and v are mutual translations, then $p_{(u,v)}$ tends to a relatively high probability, which we will call λ^+ . If u and v are not mutual translations, then $p_{(u,v)}$ tends to a very low probability, which we will call λ^- . λ^+ and λ^- correspond to the two peaks in the frequency distribution of $k_{(u,v)}/n_{(u,v)}$ in Figure 2. The two parameters can also be interpreted as the percentage of true and false positives. If the translation in the bitext is consistent and the

model is accurate, then λ^+ should be near 1 and λ^- should be near 0.

To find the most probable values of the hidden model parameters λ^+ and λ^- , we adopt the standard method of maximum likelihood estimation, and find the values that maximize the probability of the link frequency distributions. The one-to-one assumption implies independence between different link types, so that

$$\Pr(\text{links}|\text{model}) = \prod_{\mathbf{u}, \mathbf{v}} \Pr(k_{(\mathbf{u}, \mathbf{v})} | n_{(\mathbf{u}, \mathbf{v})}, \lambda^+, \lambda^-). \quad (1)$$

The factors on the right-hand side of Equation 1 can be written explicitly with the help of a mixture coefficient. Let τ be the probability that an arbitrary co-occurring pair of word types are mutual translations. Let $B(k|n, p)$ denote the probability that k links are observed out of n co-occurrences, where k has a binomial distribution with parameters n and p . Then the probability that \mathbf{u} and \mathbf{v} are linked $k_{(\mathbf{u}, \mathbf{v})}$ times out of $n_{(\mathbf{u}, \mathbf{v})}$ co-occurrences is a mixture of two binomials:

$$\begin{aligned} \Pr(k_{(\mathbf{u}, \mathbf{v})} | n_{(\mathbf{u}, \mathbf{v})}, \lambda^+, \lambda^-) &= \\ &= \tau B(k_{(\mathbf{u}, \mathbf{v})} | n_{(\mathbf{u}, \mathbf{v})}, \lambda^+) \\ &+ (1 - \tau) B(k_{(\mathbf{u}, \mathbf{v})} | n_{(\mathbf{u}, \mathbf{v})}, \lambda^-). \end{aligned} \quad (2)$$

One more variable allows us to express τ in terms of λ^+ and λ^- : Let λ be the probability that an arbitrary co-occurring pair of word tokens will be linked, regardless of whether they are mutual translations. Since τ is constant over all word types, it also represents the probability that an arbitrary co-occurring pair of *tokens* are mutual translations. Therefore,

$$\lambda = \tau \lambda^+ + (1 - \tau) \lambda^-. \quad (3)$$

λ can also be estimated empirically. Let K be the total number of links in the bitext and let N be the total number of co-occurring word token pairs: $K = \sum_{(\mathbf{u}, \mathbf{v})} k_{(\mathbf{u}, \mathbf{v})}$, $N = \sum_{(\mathbf{u}, \mathbf{v})} n_{(\mathbf{u}, \mathbf{v})}$. By definition,

$$\lambda = K/N. \quad (4)$$

Equating the right-hand sides of Equations (3) and (4) and rearranging the terms, we get:

$$\tau = \frac{K/N - \lambda^-}{\lambda^+ - \lambda^-}. \quad (5)$$

Since τ is now a function of λ^+ and λ^- , only the latter two variables represent degrees of freedom in the model.

The probability function expressed by Equations 1 and 2 has many local maxima. In practice, these




Figure 4: $\Pr(\text{links}|\text{model})$ has only one global maximum in the region of interest.

local maxima are like pebbles on a mountain, invisible at low resolution. We computed Equation 1 over various combinations of λ^+ and λ^- after the model's first iteration over 300000 aligned sentence pairs from the Canadian Hansard bitext. Figure 4 shows that the region of interest in the parameter space, where $1 > \lambda^+ > \lambda > \lambda^- > 0$, has only one clearly visible global maximum. This global maximum can be found by standard hill-climbing methods, as long as the step size is large enough to avoid getting stuck on the pebbles.

Given estimates for λ^+ and λ^- , we can compute $B(k_{\mathbf{u}, \mathbf{v}} | n_{\mathbf{u}, \mathbf{v}}, \lambda^+)$ and $B(k_{\mathbf{u}, \mathbf{v}} | n_{\mathbf{u}, \mathbf{v}}, \lambda^-)$. These are probabilities that $k_{(\mathbf{u}, \mathbf{v})}$ links were generated by an algorithm that generates correct links and by an algorithm that generates incorrect links, respectively, out of $n_{(\mathbf{u}, \mathbf{v})}$ co-occurrences. The ratio of these probabilities is the likelihood ratio in favor of \mathbf{u} and \mathbf{v} being mutual translations, for all \mathbf{u} and \mathbf{v} :

$$L(\mathbf{u}, \mathbf{v}) = \frac{B(k_{\mathbf{u}, \mathbf{v}} | n_{\mathbf{u}, \mathbf{v}}, \lambda^+)}{B(k_{\mathbf{u}, \mathbf{v}} | n_{\mathbf{u}, \mathbf{v}}, \lambda^-)}. \quad (6)$$

4 Class-Based Word-to-Word Models

In the basic word-to-word model, the hidden parameters λ^+ and λ^- depend only on the distributions of link frequencies generated by the competitive linking algorithm. More accurate models can be induced by taking into account various features of the linked tokens. For example, frequent words are translated less consistently than rare words (Melamed, 1997). To account for this difference, we can estimate separate values of λ^+ and λ^- for different ranges of $n_{(\mathbf{u}, \mathbf{v})}$. Similarly, the hidden parameters can be conditioned on the linked parts of speech. Word order can be taken into account by conditioning the hidden parameters on the relative positions of linked word tokens in their respective sentences. Just as easily, we can model links that coincide with entries in a pre-existing translation lexicon separately

from those that do not. This method of incorporating dictionary information seems simpler than the method proposed by Brown et al. for their models (Brown et al., 1993b). When the hidden parameters are conditioned on different link classes, the estimation method does not change; it is just repeated for each link class.

5 Evaluation

A word-to-word model of translational equivalence can be evaluated either over types or over tokens. It is impossible to replicate the experiments used to evaluate other translation models in the literature, because neither the models nor the programs that induce them are generally available. For each kind of evaluation, we have found one case where we can come close.

We induced a two-class word-to-word model of translational equivalence from 13 million words of the Canadian Hansards, aligned using the method in (Gale & Church, 1991). One class represented content-word links and the other represented function-word links⁴. Link types with negative log-likelihood were discarded after each iteration. Both classes' parameters converged after six iterations. The value of class-based models was demonstrated by the differences between the hidden parameters for the two classes. (λ^+ , λ^-) converged at (.78,.00016) for content-class links and at (.43,.000094) for function-class links.

5.1 Link Types

The most direct way to evaluate the link types in a word-level model of translational equivalence is to treat each link type as a candidate translation lexicon entry, and to measure precision and recall. This evaluation criterion carries much practical import, because many of the applications mentioned in Section 1 depend on accurate broad-coverage translation lexicons. Machine readable bilingual dictionaries, even when they are available, have only limited coverage and rarely include domain-specific terms (Resnik & Melamed, 1997).

We define the recall of a word-to-word translation model as the fraction of the bitext vocabulary represented in the model. Translation model precision is a more thorny issue, because people disagree about the degree to which context should play a role in judgements of translational equivalence. We hand-evaluated the precision of the link types in our model in the context of the bitext from which the model

⁴Since function words can be identified by table lookup, no POS-tagger was involved.

was induced, using a simple bilingual concordancer. A link type (u, v) was considered correct if u and v ever co-occurred as direct translations of each other. Where the one-to-one assumption failed, but a link type captured part of a correct translation, it was judged "incomplete." Whether incomplete links are correct or incorrect depends on the application.




Figure 5: Link type precision with 95% confidence intervals at varying levels of recall.

We evaluated five random samples of 100 link types each at three levels of recall. For our bitext, recall of 36%, 46% and 90% corresponded to translation lexicons containing 32274, 43075 and 88633 words, respectively. Figure 5 shows the precision of the model with 95% confidence intervals. The upper curve represents precision when incomplete links are considered correct, and the lower when they are considered incorrect. On the former metric, our model can generate translation lexicons with precision and recall both exceeding 90%, as well as dictionary-sized translation lexicons that are over 99% correct.

Though some have tried, it is not clear how to extract such accurate lexicons from other published translation models. Part of the difficulty stems from the implicit assumption in other models that each word has only one sense. Each word is assigned the same unit of probability mass, which the model distributes over all candidate translations. The correct translations of a word that has several correct translations will be assigned a lower probability than the correct translation of a word that has only one correct translation. This imbalance foils thresholding strategies, clever as they might be (Gale & Church, 1991; Wu & Xia, 1994; Chen, 1996). The likelihoods in the word-to-word model remain unnormalized, so they do not compete.

The word-to-word model maintains high precision even given much less training data. Resnik & Melamed (1997) report that the model produced

translation lexicons with 94% precision and 30% recall, when trained on French/English software manuals totaling about 400,000 words. The model was also used to induce a translation lexicon from a 6200-word corpus of French/English weather reports. Nasr (1997) reported that the translation lexicon that our model induced from this tiny bitext accounted for 30% of the word types with precision between 84% and 90%. Recall drops when there is less training data, because the model refuses to make predictions that it cannot make with confidence. For many applications, this is the desired behavior.

5.2 Link Tokens

type of error	errors made by IBM Model 2	errors made by our model
wrong link	32	7
missing link	12	36
partial link	7	10
class conflict	-	5
tokenization	3	2
paraphrase	39	36
TOTAL	93	96

Table 1: *Erroneous link tokens generated by two translation models.*

The most detailed evaluation of link tokens to date was performed by (Macklovitch & Hannan, 1996), who trained Brown et al.'s Model 2 on 74 million words of the Canadian Hansards. These authors kindly provided us with the links generated by that model in 51 aligned sentences from a held-out test set. We generated links in the same 51 sentences using our two-class word-to-word model, and manually evaluated the content-word links from both models. The IBM models are directional; i.e. they posit the English words that gave rise to each French word, but ignore the distribution of the English words. Therefore, we ignored English words that were linked to nothing.

The errors are classified in Table 1. The “wrong link” and “missing link” error categories should be self-explanatory. “Partial links” are those where one French word resulted from multiple English words, but the model only links the French word to one of its English sources. “Class conflict” errors resulted from our model’s refusal to link content words with function words. Usually, this is the desired behavior, but words like English auxiliary verbs are sometimes used as content words, giving rise to content words in French. Such errors could be overcome by a model that classifies each word token, for example using a

part-of-speech tagger, instead of assigning the same class to all tokens of a given type. The bitext preprocessor for our word-to-word model split hyphenated words, but Macklovitch & Hannan’s preprocessor did not. In some cases, hyphenated words were easier to link correctly; in other cases they were more difficult. Both models made some errors because of this tokenization problem, albeit in different places. The “paraphrase” category covers all link errors that resulted from paraphrases in the translation. Neither IBM’s Model 2 nor our model is capable of linking multi-word sequences to multi-word sequences, and this was the biggest source of error for both models.

The test sample contained only about 400 content words⁵, and the links for both models were evaluated post-hoc by only one evaluator. Nevertheless, it appears that our word-to-word model with only two link classes does not perform any worse than IBM’s Model 2, even though the word-to-word model was trained on less than one fifth the amount of data that was used to train the IBM model. Since it doesn’t store indirect associations, our word-to-word model contained an average of 4.5 French words for every English word. Such a compact model requires relatively little computational effort to induce and to apply.




Figure 6: *An example of the different sorts of errors made by the word-to-word model and the IBM Model 2. Solid lines are links made by both models; dashes lines are links made by the IBM model only. Only content-class links are shown. Neither model makes the correct links (déchainés, screaming) and (démontée, dangerous).*

⁵The exact number depends on the tokenization method.

In addition to the quantitative differences between the word-to-word model and the IBM model, there is an important qualitative difference, illustrated in Figure 6. As shown in Table 1, the most common kind of error for the word-to-word model was a missing link, whereas the most common error for IBM's Model 2 was a wrong link. Missing links are more informative: they indicate where the model has failed. The level at which the model trusts its own judgement can be varied directly by changing the likelihood cutoff in Step 1 of the competitive linking algorithm. Each application of the word-to-word model can choose its own balance between link token precision and recall. An application that calls on the word-to-word model to link words in a bitext could treat unlinked words differently from linked words, and avoid basing subsequent decisions on uncertain inputs. It is not clear how the precision/recall trade-off can be controlled in the IBM models.

One advantage that Brown et al.'s Model 1 has over our word-to-word model is that their objective function has no local maxima. By using the EM algorithm (Dempster et al., 1977), they can guarantee convergence towards the globally optimum parameter set. In contrast, the dynamic nature of the competitive linking algorithm changes the $\Pr(\text{data}|\text{model})$ in a non-monotonic fashion. We have adopted the simple heuristic that the model "has converged" when this probability stops increasing.

6 Conclusion

Many multilingual NLP applications need to translate words between different languages, but cannot afford the computational expense of modeling the full range of translation phenomena. For these applications, we have designed a fast algorithm for estimating word-to-word models of translational equivalence. The estimation method uses a pair of hidden parameters to measure the model's uncertainty, and avoids making decisions that it's not likely to make correctly. The hidden parameters can be conditioned on information extrinsic to the model, providing an easy way to integrate pre-existing knowledge.

So far we have only implemented a two-class model, to exploit the differences in translation consistency between content words and function words. This relatively simple two-class model linked word tokens in parallel texts as accurately as other translation models in the literature, despite being trained on only one fifth as much data. Unlike other translation models, the word-to-word model can automat-

ically produce dictionary-sized translation lexicons, and it can do so with over 99% accuracy.

Even better accuracy can be achieved with a more fine-grained link class structure. Promising features for classification include part of speech, frequency of co-occurrence, relative word position, and translational entropy (Melamed, 1997). Another interesting extension is to broaden the definition of a "word" to include multi-word lexical units (Smadja, 1992). If such units can be identified *a priori*, their translations can be estimated without modifying the word-to-word model. In this manner, the model can account for a wider range of translation phenomena.

Acknowledgements

The French/English software manuals were provided by Gary Adams of Sun Microsystems Laboratories. The weather bitext was prepared at the University of Montreal, under the direction of Richard Kit-tredge. Thanks to Alexis Nasr for hand-evaluating the weather translation lexicon. Thanks also to Mike Collins, George Foster, Mitch Marcus, Lyle Ungar, and three anonymous reviewers for helpful comments. This research was supported by an equipment grant from Sun Microsystems and by ARPA Contract #N66001-94C-6043.

References

- P. F. Brown, J. Cocke, S. Della Pietra, V. Della Pietra, F. Jelinek, R. Mercer, & P. Roossin, "A Statistical Approach to Language Translation," *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, 1988.
- P. F. Brown, J. Cocke, S. Della Pietra, V. Della Pietra, F. Jelinek, R. Mercer, & P. Roossin, "A Statistical Approach to Machine Translation," *Computational Linguistics* 16(2), 1990.
- P. F. Brown, V. J. Della Pietra, S. A. Della Pietra & R. L. Mercer, "The Mathematics of Statistical Machine Translation: Parameter Estimation," *Computational Linguistics* 19(2), 1993.
- P. F. Brown, S. A. Della Pietra, V. J. Della Pietra, M. J. Goldsmith, J. Hajic, R. L. Mercer & S. Mohanty, "But Dictionaries are Data Too," *Proceedings of the ARPA HLT Workshop*, Princeton, NJ, 1993.
- R. Catizone, G. Russell & S. Warwick "Deriving Translation Data from Bilingual Texts," *Proceedings of the First International Lexical Acquisition Workshop*, Detroit, MI, 1993.

- S. Chen, *Building Probabilistic Models for Natural Language*, Ph.D. Thesis, Harvard University, 1996.
- K. W. Church & E. H. Hovy, "Good Applications for Crummy Machine Translation," *Machine Translation 8*, 1993.
- I. Dagan, K. Church, & W. Gale, "Robust Word Alignment for Machine Aided Translation," *Proceedings of the Workshop on Very Large Corpora: Academic and Industrial Perspectives*, Columbus, OH, 1993.
- A. P. Dempster, N. M. Laird & D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society 34(B)*, 1977.
- T. Dunning, "Accurate Methods for the Statistics of Surprise and Coincidence," *Computational Linguistics 19(1)*, 1993.
- P. Fung, "Compiling Bilingual Lexicon Entries from a Non-Parallel English-Chinese Corpus," *Proceedings of the Third Workshop on Very Large Corpora*, Boston, MA, 1995a.
- P. Fung, "A Pattern Matching Method for Finding Noun and Proper Noun Translations from Noisy Parallel Corpora," *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, Boston, MA, 1995b.
- W. Gale & K. W. Church, "A Program for Aligning Sentences in Bilingual Corpora" *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, CA, 1991.
- W. Gale & K. W. Church, "Identifying Word Correspondences in Parallel Texts," *Proceedings of the DARPA SNL Workshop*, 1991.
- A. Kumano & H. Hirakawa, "Building an MT Dictionary from Parallel Texts Based on Linguistic and Statistical Information," *Proceedings of the 15th International Conference on Computational Linguistics*, Kyoto, Japan, 1994.
- E. Macklovitch "Using Bi-textual Alignment for Translation Validation: The TransCheck System," *Proceedings of the 1st Conference of the Association for Machine Translation in the Americas*, Columbia, MD, 1994.
- E. Macklovitch & M.-L. Hannan, "Line 'Em Up: Advances in Alignment Technology and their Impact on Translation Support Tools," *2nd Conference of the Association for Machine Translation in the Americas*, Montreal, Canada, 1996.
- I. D. Melamed "Automatic Evaluation and Uniform Filter Cascades for Inducing N -best Translation Lexicons," *Proceedings of the Third Workshop on Very Large Corpora*, Boston, MA, 1995.
- I. D. Melamed, "A Geometric Approach to Mapping Bitext Correspondence," *Proceedings of the First Conference on Empirical Methods in Natural Language Processing*, Philadelphia, PA, 1996a.
- I. D. Melamed "Automatic Detection of Omissions in Translations," *Proceedings of the 16th International Conference on Computational Linguistics*, Copenhagen, Denmark, 1996b.
- I. D. Melamed, "Automatic Construction of Clean Broad-Coverage Translation Lexicons," *2nd Conference of the Association for Machine Translation in the Americas*, Montreal, Canada, 1996c.
- I. D. Melamed, "Measuring Semantic Entropy," *Proceedings of the SIGLEX Workshop on Tagging Text with Lexical Semantics*, Washington, DC, 1997.
- I. D. Melamed, "A Portable Algorithm for Mapping Bitext Correspondence," *Proceedings of the 35th Conference of the Association for Computational Linguistics*, Madrid, Spain, 1997. (in this volume)
- A. Melby, "A Bilingual Concordance System and its Use in Linguistic Studies," *Proceedings of the English LACUS Forum*, Columbia, SC, 1981.
- A. Nasr, personal communication, 1997.
- P. Resnik & I. D. Melamed, "Semi-Automatic Acquisition of Domain-Specific Translation Lexicons," *Proceedings of the 7th ACL Conference on Applied Natural Language Processing*, Washington, DC, 1997.
- D. W. Oard & B. J. Dorr, "A Survey of Multilingual Text Retrieval," *UMIACS TR-96-19*, University of Maryland, College Park, MD, 1996.
- F. Smadja, "How to Compile a Bilingual Collocational Lexicon Automatically," *Proceedings of the AAAI Workshop on Statistically-Based NLP Techniques*, 1992.
- D. Wu & X. Xia, "Learning an English-Chinese Lexicon from a Parallel Corpus," *Proceedings of the First Conference of the Association for Machine Translation in the Americas*, Columbia, MD, 1994.