Crosslinguistic Transfer in Automatic Verb Classification

Vivian Tsang
Computer Science
University of Toronto
vyctsang@cs.toronto.edu

Suzanne Stevenson
Computer Science
University of Toronto
suzanne@cs.toronto.edu

Paola Merlo
Linguistics Dept.
University of Geneva
merlo@lettres.unige.ch

Abstract

We investigate the use of multilingual data in the automatic classification of English verbs, and show that there is a useful transfer of information across languages. Specifically, we experiment with three lexical semantic classes of English verbs. We collect statistical features over a sample of English verbs from each of the classes, as well as over Chinese translations of those verbs. We use the English and Chinese data, alone and in combination, as training data for a machine learning algorithm whose output is an automatic verb classifier. We demonstrate that Chinese data is indeed useful in helping to classify the English verbs (at 82% accuracy), and furthermore that a multilingual combination of data outperforms the English data alone (85% accuracy). Moreover, our results using monolingual corpora show that it is not necessary to use a parallel corpus to extract the translations in order for this technique to be successful.

1 Introduction

Automatic acquisition of lexical information is critical to the creation of lexicons for wide-coverage NLP systems. Recently, a number of researchers have devised corpus-based approaches for automatically learning the lexical semantic classes of verbs (e.g., McCarthy and Korhonen (1998); Lapata and Brew (1999); Schulte im Walde (2000); Merlo and Stevenson (2001)). Such classes incorporate complex syntactic and semantic information common to a set of verbs that share a general semantic property (such as expressing a manner of motion, or a change of state) (Kipper et al., 2000). Automatic classification of verbs can thus avoid the need for expensive manual coding.

Corpus-based approaches to this problem rest on an assumption of regularity in the mapping from the semantics of a verb, to its syntactic usage (Levin, 1993). Statistical features over the syntax of a verb can be informative about its

underlying semantic classification, and can thus serve as the training data for an automatic classifier (Merlo and Stevenson, 2001). The drawback to this approach is that the features that are readily extractable from a corpus are typically only noisy or indirect indicators of the semantics—and indeed, some relevant semantic properties may simply not be expressed overtly.

Interestingly, however, languages differ in the syntactic expression of semantic properties; e.g., in English, change of state verbs are used in the same form in either of their non-causative or causative senses, but in Chinese, a causative meaning often requires the use of an overt periphrastic particle with the verb. Thus, a particular mapping between semantics and syntax may be more easily detected in one language than in another.

In our work, we exploit this crosslinguistic variation by using multilingual data to classify verbs in a single language (our test case is English). Our motivating observation is that some semantic distinctions that are difficult to detect superficially in English may manifest themselves as surface syntactic indicators in another language (we have experimented with Chinese, Italian, and German; here we focus on Chinese). Thus, we should be able to augment the set of features for English verb classes that we extract from an English corpus, with features over the translations of the English verbs in the other language(s).

Our work is guided by several hypotheses about the relation across languages between the semantics of a verb, and its expression in syntactic corpus data. First, we hypothesize that a second language can provide data that will be helpful in classifying English verbs. Note that we do not rely on there being a universal semantic classification of verbs across languages. The verbs within an English class share some general semantic property (such as change of state), which leads to a commonality of syntactic behav-

ior among them. We hypothesize that verbs in the translation set for the English class will also share this property, at least in sufficient numbers to lead to consistent syntactic behavior among the translation set as well. In support of this first hypothesis, we find that statistics over Chinese verbs do indeed perform very well, on their own, in classifying the English verbs that they are translations of.

Second, because we are relying on a general semantic overlap between verbs in the two languages, and not on a one-to-one correspondence in usage, we hypothesize that a parallel corpus is unnecessary to the success of the transfer of information across languages. We should be able to determine the translations, and extract their associated data, from a monolingual corpus. In preliminary work, Tsang and Stevenson (2001) showed the usefulness of a parallel corpus in this kind of approach (on a smaller number of verbs and classes than investigated here). Here, our Chinese translations and features are determined from a (larger) monolingual Chinese corpus, and in fact outperform the earlier results obtained using a parallel corpus.

Our final hypothesis is that combinations of features across the two languages will be even more helpful than either set of features on its own. If the data from different languages really represent the differing syntactic expression of a common underlying semantics, then the multilingual combination of data will provide differing views of the same classification to the machine learning algorithm, increasing the usefulness of its training data. Again, our hypothesis is supported, as we find that, in almost all of our experiments, a combination of English and Chinese features outperforms either feature set used monolingually.

In the remainder of the paper, we describe our sample English verb classes and features, the determination of Chinese features with which to augment these, and the machine learning experiments on both monolingual and multilingual combinations of the features.

2 English Verb Classes and Features

Following the approach of Merlo and Stevenson (2001), henceforward MS01, we focus on a lexical semantic classification of verbs based on their argument structure. By argument structure, we mean the thematic roles assigned by a verb—

such as Agent or Theme—and their mapping to syntactic positions. This provides a similar type of semantic classification to that of Levin (1993), although at a coarser level, since it relies solely on general participant role information and not on fine-grained semantics.

Specifically, we investigated three English classes whose verbs can appear both transitively and intransitively, but differ in argument structure. While our definition of classes is broader than the fine-grained classification developed by Levin (1993), argument structure classes generally correspond to her broader groupings of verbs (such as, e.g., class 45 instead of 45.1 or 45.2). In our case, we look at: manner of motion verbs (Levin’s class 51), change of state verbs (Levin’s 45), and verbs of creation and transformation (Levin’s 26).

The manner of motion and change of state verbs participate in a causative transitive, which inserts a causal agent into the argument structure of the single-argument intransitive form. The creation and transformation verbs can simply drop their optional object. The following sentence pairs exemplify the relation between the transitive and intransitive forms for each class:

Manner of motion:

The lion jumped through the hoop.

The trainer jumped the lion through the hoop.

Change of State:

The butter melted in the pan.

The cook melted the butter in the pan.

Creation/Transformation:

The contractor built the houses last summer.

The contractor built all summer.¹

Table 1 shows that each class is uniquely distinguished by the pattern of thematic roles assigned within the constructions above.

Classes	Transitive		Intrans
	Subj	Obj	Subj
MannerOfMotion	Caus	Ag	Ag
ChangeOfState	Caus	Th	Th
Creation/Trans	Ag	Th	Ag

Table 1: Thematic Roles by Class.

Ag=Agent, Th=Theme, Caus=Causal Agent

¹The progressive, as in *The contractor was building all summer*, may be more natural for some verbs in this usage.

In the MS01 proposal, the thematic properties of these classes were analysed to determine features that could discriminate the classes within an automatic classification system.² The result was a set of 5 numeric indicators encoding summary statistics over the usage of each verb across the Wall Street Journal (WSJ, 65M words); all were normalized frequency counts over tagged or parsed text, which had no semantic annotation. The statistical features were shown to approximate the verbs’ thematic relations, either directly or indirectly. The features are: animacy of subject—ANIM—indicating agentivity; transitive use, calculated in several variants—TRANS, PASS (passive use), VBN (passive participle tag)—indicating degree of markedness of the transitive argument structure; and use in a causative transitive—CAUS—indicating the use of a causal agent. These features successfully contributed to the classification of English verbs in MS01’s monolingual experiments, achieving an accuracy of almost 70% in a task with a 34% baseline. We adopted these same features as the starting point of our multilingual work.

3 Chinese Features

Given the English verb classes and features described above, the next step is to determine features over the Chinese translations of the verbs that could complement the English features. Recall our observation that some semantic properties of verbs may be expressed more overtly in one language (such as Chinese) than in another (such as English). Not only are such overt indicators easy to extract from a corpus by automatic means, they have the potential to enrich the existing English features, providing more information to the learning algorithm regarding the underlying thematic distinctions between the classes.

In this context, the following features were investigated. In all cases, the features are calculated as the normalized frequency of occurrence of the syntactic property described.

Chinese POS-tags for Verbs We used the POS-tagger provided by the Chinese Knowledge

²In on-going work, Joanis and Stevenson (In preparation) have explored a general feature space for capturing verb class distinctions, that eliminates the need for manually determining the features for discriminating particular classes. Preliminary experiments have achieved very good accuracies on a number of different classes.

Information Processing Group (CKIP) to automatically assign one of 15 verb tags to each verb. Each tag incorporates both subcategorization information as well as the stative/active distinction. (What is considered “stative” in Chinese is quite similar to what can be adjectivized in English—in this case, the change of state verbs.) This feature thus indicates degree of transitivity, analogously to the English transitivity features, as well as additional semantic information lacking in any of the English features.

Passive Particles In Chinese, a passive construction is indicated by a passive particle preceding the main verb. For example, *This store is closed by the owner* can be translated as *Zhe ge (this) shang dian (store) bei (passive particle) dong zhu (owner) guan bi (closed)*. Passive particles are similar to the English passive feature in indicating transitivity, but differ in their ease of detection compared to passive in English.

Periphrastic (Causative) Particles In Chinese, some causative sentences use an external (periphrastic) particle to indicate that the subject is the causal agent of the event specified by the verb. For example, one possible translation for *I cracked an egg* is *Wo (I) jiang (made, periphrastic particle) dan (egg) da lan (crack)*. This feature is analogous to the English causative feature, though (as with the passive construction) the particle expresses causativity more overtly in Chinese.

Morpheme Information We also investigated other features that captured statistics over the precise morphemic constitution of the Chinese translations (such as compound V-N or V-V verbs). Since these features proved to not be highly useful in classification, we will not discuss them further.

The four general types of features we describe above lead to 28 Chinese features in total, although in practice a number of the verb tag features are unused, since they are not applied to the verbs in our translation set. We refer to the Chinese features as follows: CKIP for the set of verb tag features, C-PASS for the passive particle feature, and C-CAUS for the causative particle feature. The Chinese features can be used alone or in combination with the 5 English features proposed by MS01.

4 Materials and Method

We chose 20 English verbs per class, and extracted their features from the British National Corpus (BNC, 100M words), which had been POS-tagged (Brill, 1995) and chunked (Abney, 1996). All counts were collected based on the combined output of the tagger and the chunker, except VBN which relied solely on the tagger’s output. The value of an English feature for a verb is the normalized frequency of the counts.

To collect the Chinese data, we need to determine our translation sets first. To find the translations, we used a portion of the Mandarin Chinese News Text (MNews, People’s Daily and Xinhua newswire sections, approximately 165M characters). We tagged the corpus using the CKIP tagger mentioned earlier, then automatically extracted all Chinese compounds with a verb POS-tag, resulting in a total of 36,323 unique verb instances. We manually selected those that are translations of the 60 English verbs in the appropriate semantic meaning, i.e., manner of motion, change of state, and creation/transformation.³ Note that since we are not classifying the Chinese verbs, we can use multiple translations per English verb, yielding more data; on average, each English verb has 6.5 translations.

The Chinese features are calculated as follows. The required counts are collected partly automatically (CKIP, C-PASS, C-CAUS) and partly by hand (morpheme combinations). The value of a Chinese feature for an English verb is the normalized frequency of occurrence of the feature across all occurrences of all Chinese verbs in the translation set of the English verb. That is, if $C_1, \dots, C_i$ are translations of the English verb E_j , then the value of Chinese feature c_k for E_j is the normalized frequency of counts across all occurrences of $C_1, \dots, C_i$.

The data for our machine learning experiments consists of a vector of the English and Chinese features for each English verb:

Template: [verb, $e_1, \dots, e_5, c_1, \dots, c_{28}$, class]

Example: [change, 0.04, $\dots$, 1, change-of-state]

where $e_1, \dots, e_5$ and $c_1, \dots, c_{28}$ are the 5 possible En-

glish features and the 28 possible Chinese features, respectively, for a total of 33 features.

We use the resulting vectors as the training data for the C5.0 machine learning system, which uses a decision tree induction algorithm (<http://www.rulequest.com>). We used a 10-fold cross-validation methodology (repeated 50 times) for our experiments.⁴ The cross-validation experiments train on a large number of random subsets of the data, for which we report average accuracy and standard error. To evaluate the contribution of different features to learning, and find the best feature combination(s), we varied the precise set of features used in each experiment. We analysed the performance of subsets of monolingual features, and the performance of combinations of features across the two languages. We also performed experiments on each pair of verb classes (three extra sets of experiments), in order to evaluate which feature combinations are most effective in distinguishing each pair of classes.

5 Experimental Results

We report here the key results of our cross-validation experiments. Recall that we have 20 English verbs per class. Hence, the baseline (chance) accuracy is 33.3% (20/60) for the 3-way experiments, and 50% (20/40) for the pairwise experiments. Although the theoretical maximum accuracy is 100%, it is worth noting that, for their 3-way verb classification task on a similar set of verbs, MS01 experimentally determined a best performance of 87% among a group of human experts, indicating that a more realistic upper-bound for the 3-way automatic classification task falls well below 100%.

Before turning to a detailed analysis of the results, it is worth briefly reviewing our guiding hypotheses: that Chinese data from a monolingual corpus could be helpful in English verb classification, and that a combination of English and Chinese features should be most useful. Tables 2 to 5 each report three results—the performance on the best subset of English-only features, on the best subset of Chinese-only features, and on the best multilingual subset of features. This allows us to analyse the Chinese-only performance, and

³Clearly, verbs can be ambiguous, and our corpora are not sense-tagged. As in MS01, we assume that the statistical features will reflect the predominant sense in the corpus.

⁴A 10-fold cross-validation experiment divides the data into ten parts and runs 10 times, each time training on a different 90% of the data and testing on the remaining 10%.

to compare the best monolingual performance, in either language, to the best multilingual combination.⁵

Features	%Acc.	%SE
Best English: ANIM, TRANS	67.6	0.3
Best Chinese: CKIP	82.1	0.1
Best multi: ANIM, PASS, CKIP, C-PASS	85.2	0.3

Table 2: Three-way classification accuracy using 10-fold cross-validation, 50 repeats

Features	%Acc.	%SE
Best English: All	88.1	0.4
Best Chinese: CKIP	92.4	0.1
Best multi: Any Eng, CKIP	92.6	0.2

Table 3: 2-way classification (manner of motion, change of state) accuracy using 10-fold cross-validation, 50 repeats

First consider the results of our 3-way classification experiment, shown in Table 2. As predicted, the Chinese features perform very well alone, at an accuracy of 82.1% using the CKIP features. Indeed, the Chinese features outperform the best English features of ANIM and TRANS, which attained 67.6%. Clearly, features from a second language can be very useful—even more useful than English features—in English verb classification. Additionally, we see that the combination of English and Chinese features consisting of ANIM, PASS, CKIP, and C-PASS, achieves the highest performance of 85.2%—a very good result on a task with a 33.3% baseline. Thus, a multilingual combination of features appears to yield more information to the automatic classifier, as expected. Both of our hypotheses, then, receive strong support from these results.

In order to gain some insight into why particular features were helpful in the 3-way classification task, we turn to the results of our pairwise experiments, to determine which features are most useful in distinguishing each pair of classes.

⁵Note that in all cases except one (where the reported accuracies of two experiments are 92.6% and 92.4%), the difference between each pair of accuracies in a table is significant at $p < 0.05$, using a one-way ANOVA with Tukey-Kramer post-tests.

Features	%Acc.	%SE
Best English: VBN	82.5	0.0
Best Chinese: CKIP	90.1	0.2
Best multi: PASS, CKIP	93.7	0.3

Table 4: 2-way classification (manner of motion, creation/transformation) accuracy using 10-fold cross-validation, 50 repeats

Features	%Acc.	%SE
Best English: ANIM	80.3	0.2
Best Chinese: CKIP	81.8	0.2
Best multi: All Eng, CKIP OR ANIM, CKIP, C-PASS	86.7	0.3

Table 5: 2-way classification (change of state, creation/transformation) accuracy using 10-fold cross-validation, 50 repeats

(Recall that in these experiments, the baseline accuracy is higher—50% instead of 33%.)

Table 3 shows the results of the experiments on the manner of motion and change of state classes. Here we find that the CKIP features perform best overall (92.4%), with no additional advantage to combination with English features. (The difference between 92.4% and 92.6% is not statistically significant.) Table 4 gives the results for the manner of motion and creation/transformation classes. In this case, the best performance is the multilingual combination of CKIP and PASS (93.7%). Finally, Table 5 shows the results of the experiments on the change of state and creation/transformation classes. Again, the best result is with a multilingual combination, in this case either CKIP with all the English features, or CKIP with C-PASS and ANIM (86.7%).

On the one hand, CKIP appears to be helpful in distinguishing all three pairs of classes; on the other hand, it participates in a different combination of features in the best result for each pairwise comparison. Let’s first consider the usefulness of CKIP. It turns out that several of the verb tags are directly relevant to our classes. In comparing the feature values for these tags across the three classes, we find that VC, the transitive/active tag, makes a 3-way distinction that mirrors the transitivity distinction

in English (creation/transformation most transitive, and manner of motion least); VA, the intransitive/active tag, distinguishes manner of motion (activity verbs that are primarily intransitive) from the other two classes; and VH/VHC, the intransitive/stative tags, distinguish change of state (the only stative class in our group) from the other two classes. It is not surprising, then, that the feature combination including these tags helps distinguish all three pairs of classes.

But while CKIP is sufficient for high accuracy on the manner of motion/change of state distinction, it must be combined with different features for the other pairs of classes. For manner of motion and creation/transformation verbs, the PASS feature is best in combination with CKIP. PASS is another indicator of transitivity, so these two features appear to complement each other in providing different descriptions across the languages of similar properties. Interestingly, for the change of state and creation/transformation distinction, again the passive is a helpful feature, but this time it is the passive in Chinese, C-PASS. So different views (CKIP and C-PASS) of the same property (transitivity) can be useful within a language as well. In addition, ANIM is useful in combination for the change of state and creation/transformation distinction. ANIM indicates a different salient property of these classes: the creation/transformation verbs are more likely to have an agentive (and therefore animate) subject.

In returning to our 3-way results, it is now possible to understand the particular combination of features that performs the best, ANIM, PASS, CKIP, and C-PASS. Each of these features participates in the best combination for one or more of the three pairwise experiments. Thus, the best 3-way performance is achieved by taking a union of the features that perform best in the 2-way experiments. This is a useful outcome, since it enables us to better understand the differing, but stable, contributions of the features to the results.

To summarize our main results, with the exception of the two-way experiment between manner of motion and change of state verbs, a multilingual combination of features consistently outperforms either set of monolingual features. In the one case where they don't, it is the Chinese features that perform best overall. Indeed, in all

cases, the best Chinese features alone outperform the best English features alone. These results provide strong support for our motivating hypotheses—that Chinese features (even those extracted from a monolingual corpus) will be useful in English verb classification, and even more useful in combination with English features.

Specifically, English animacy and transitivity, Chinese POS-tags, and the passive feature in both languages distinguish the pairs of classes, and all three classes, quite well. It is worth noting that the performance of a particular feature in one language is an indicator of the performance of the related feature in another language. For example, both passive features do not perform well alone, but perform well in combination. On the other hand, the causative feature from neither language performs well, alone or in combination. This evidence indicates that there are syntactic/semantic properties that hold across languages, supporting crosslinguistic transfer in verb classification.

6 Related Work

Multilingual resources are widely used in several areas of NLP. The key is to exploit the underlying syntactic and/or semantic commonalities between languages. For example, Ide (2000) and Resnik and Yarowsky (1999) used parallel corpora for lexicalizing some fine-grained English senses. Yarowsky et al. (2001) examined the transferability of syntactic information using parallel corpora as well.

However, our multilingual approach does not rest on the use of parallel corpora, and in that sense is perhaps closer to the work of Dagan and Itai (1994), which used statistical data from a monolingual corpus to aid in WSD in a different language. We have also taken inspiration from work on Second Language Acquisition, in which “transfer” of knowledge from a first language to learning a second has been shown to occur in the acquisition of verb class knowledge (e.g., Helms-Park (2001); Inagaki (2001); Montrul (2001)). Finally, our work has further connections to the machine translation and lexical acquisition work of Dorr and colleagues (e.g., Dorr (1993)), which is founded on the notion of underlying semantic commonalities among verbs as the key to crosslinguistic mappings.

7 Conclusions

In this paper, we have presented evidence that there is a useful transfer of information from Chinese to English for the lexical semantic classification of verbs. We find that a classifier for English verbs that is trained either on Chinese-only features, or on both English and Chinese, reaches accuracies of 82% and 85% respectively, in a 3-way classification task with a 33% baseline. Indeed, using Chinese features alone or in combination outperforms a classifier trained on English-only data, which attains an accuracy of only 68% in the same 3-way task. These results are based on counts collected from monolingual corpora, confirming our hypothesis that a parallel corpus from which to draw the translation data is unnecessary. We conclude that successful crosslinguistic transfer is grounded in the underlying semantic similarities in the argument structure of verbs. On-going work in other languages confirms these numerical results and their underlying hypotheses. We find that both Italian and German features perform better than English ones at a similar classification of English verbs (with accuracies of 85% and 90%, respectively, in a 2-way classification task).

Acknowledgements

We gratefully acknowledge the financial support of the NSERC of Canada, and the NSF of Switzerland. We thank Eric Joanis for his help in the extraction of the English data. The Italian and German experiments referred to are in collaboration with Gianluca Allaria (Geneva) and Paul McCabe (Toronto), respectively.

Appendix

Manner of motion verbs: *crawl, float, fly, glide, hurry, jump, leap, march, parade, race, ride, sail, scurry, skate, ski, skip, swim, vault, walk, wander*.

Change of state verbs: *burn, change, close, collapse, compress, cool, crack, decrease, dissolve, divide, drain, expand, flood, fold, freeze, increase, melt, soak, solidify, stabilize*.

Creation and transformation verbs: *build, carve, chant, choreograph, compose, cut, dance, direct, draw, hammer, knit, perform, play, produce, recite, sculpt, sew, sketch, weave, write*.

References

Steven Abney. 1996. Partial parsing via finite-state cascades. In John Carroll, editor, *Proceedings of the Workshop on Robust Parsing at the 8th Summer School on Logic, Language and Information*, pages 8–15, University of Sussex.

Eric Brill. 1995. Transformation-based error-driven learning and natural language processing: A case study in part-of-speech tagging. *Computational Linguistics*, 21(4):543–565.

Ido Dagan and Alon Itai. 1994. Word sense disambiguation using a second language monolingual corpus. *Computational Linguistics*, 20(4):563–596.

Bonnie Dorr. 1993. *Machine Translation: A View from the Lexicon*. MIT Press.

Rena Helms-Park. 2001. Evidence of lexical transfer in learner syntax – the acquisition of English causatives by speakers of Hindi-Urdu and Vietnamese. *Studies in Second Language Acquisition*, 23(1):71–102.

Nancy Ide. 2000. Cross-lingual sense determination: Can it work? *Computers and the Humanities*, 34:223–234.

Shunji Inagaki. 2001. Motion verbs with goal PPs in the L2 acquisition of English and Japanese. *Studies in Second Language Acquisition*, 23(2):153–170.

Eric Joanis and Suzanne Stevenson. In preparation. A general feature space for automatic verb classification. Manuscript, Univ. of Toronto.

Karin Kipper, Hoa Trang Dang, and Martha Palmer. 2000. Class-based construction of a verb lexicon. In *Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, Austin TX.

Maria Lapata and Chris Brew. 1999. Using subcategorization to resolve verb class ambiguity. In *Proceedings of Joint SIGDAT Conference on Empirical Methods in Natural Language Processing (EMNLP) and Very Large Corpora*, pages 266–274, College Park, MD.

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.

Diana McCarthy and Anna-Leena Korhonen. 1998. Detecting verbal participation in diathesis alternations. In *Proceedings of the 36th Annual Meeting of the ACL and the 17th International Conference on Computational Linguistics (COLING-ACL 1998)*, pages 1493–1495, Montreal, Canada.

Paola Merlo and Suzanne Stevenson. 2001. Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics*, 27(3):393–408.

Silvina Montrul. 2001. Agentive verbs of manner of motion in Spanish and English as second languages. *Studies in Second Language Acquisition*, 23(2):171–206.

Philip Resnik and David Yarowsky. 1999. Distinguishing systems and distinguishing senses: New evaluation methods for word sense disambiguation. *Natural Language Engineering*, 5(2):113–133.

Sabine Schulte im Walde. 2000. Clustering verbs semantically according to their alternation behaviour. In *Proceedings of the 18th International Conference on Computational Linguistics (COLING 2000)*, pages 747–753, Saarbrücken, Germany.

Vivian Tsang and Suzanne Stevenson. 2001. Automatic verb classification using multilingual resources. In *Proceedings of Fifth Computational Natural Language Learning Workshop*, pages 30–37.

David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. Inducing multilingual text analysis tools via robust projection across aligned corpora. In *Proceedings of HLT 2001*.