

# Resolving Out-of-Vocabulary Words with Bilingual Embeddings in Machine Translation

Pranava Swaroop Madhyastha    Cristina España-Bonet

TALP Research Center

Univesitat Politècnica de Catalunya

Jordi Girona, 1-3, 08034 Barcelona, Spain

{pranava,cristinae}@cs.upc.edu

## Abstract

Out-of-vocabulary words account for a large proportion of errors in machine translation systems, especially when the system is used on a different domain than the one where it was trained. In order to alleviate the problem, we propose to use a log-bilinear softmax-based model for vocabulary expansion, such that given an out-of-vocabulary source word, the model generates a probabilistic list of possible translations in the target language. Our model uses only word embeddings trained on significantly large unlabelled monolingual corpora and trains over a fairly small, word-to-word bilingual dictionary. We input this probabilistic list into a standard phrase-based statistical machine translation system and obtain consistent improvements in translation quality on the English–Spanish language pair. Especially, we get an improvement of 3.9 BLEU points when tested over an out-of-domain testset.

## 1 Introduction

Data-driven machine translation systems are able to translate words that have been seen in the training corpora, however translating unseen words is still a major challenge for even the best performing systems. In general, the amount of parallel data is finite (and sometimes scarce) which results in word types like named entities, domain specific content words, or infrequent terms to be absent in the training parallel corpora. This lack of information can potentially result in incomplete or erroneous translations.

This area has been actively studied in the field of machine translation (MT) (Habash, 2008; Daumé III and Jagarlamudi, 2011;

Marton et al., 2009;

Dou and Knight, 2012;

Irvine and Callison-Burch, 2013).

Lexicon based resources have been used for resolving unseen content words by exploiting a combination of monolingual and bilingual resources (Rapp, 1999; Callison-Burch et al., ; Saluja et al., 2014; Zhao et al., 2015). In this context, distributed word representations, or word embeddings (WE), have been recently applied to resolve unseen word related problems (Mikolov et al., 2013b; Zou et al., 2013). In general, word representations capture rich linguistic relationships. Several works (Gouws et al., 2015; Wu et al., 2014) try to use WE to improve MT systems. However, very few approaches use them directly to resolve the out-of-vocabulary (OOV) problem for MT systems.

Rapp, 1999;

Our work is inspired by the recent advances in applications of word embeddings to the task of vocabulary expansion in the context of statistical machine translation (SMT). In this work, we introduce a principled method to obtain a probabilistic distribution of words in the target language for a given source word. We do this by using WEs in both languages and learning a log-bilinear softmax model that is trained using a relatively small bilingual lexicon (the seed lexicon) to obtain a probabilistic distribution of words. Then, we integrate the generated distribution of target words for every unseen source word into a standard SMT system.

The rest of the paper is organised as follows. In the next section we briefly describe some previous related work. Section 3 presents the log-bilinear softmax model, and its integration into an SMT system and the SMT experiments sre analysed in Section 4. Finally, in Section 5, we draw our conclusions and sketch some future work.

## 2 Background and Motivation

There are several strands of related research that try to alleviate the effect of unseen words in translation. Previous research suggests that a significantly large number of named entities can be handled by using simple pre/post-processing, like transliteration methods (Hermjakob et al., 2008; Al-Onaizan and Knight, 2002). However, a change in domain results in a significant increase in the number of unseen words. These unseen words might include a significant proportion of regular domain-specific content words.

Our focus in this paper is to resolve unseen content words by using continuous word embeddings on both the languages and a small seed lexicon to map the embeddings. To this extent, our work is similar to Ishiwatari et al. (2016) where the authors map distributional representations using a linear regression method similar to Mikolov et al. (2013b) and insert a new feature based on cosine similarity metric into the MT system. In our work, we use a principled method to obtain a probabilistic conditional distribution of words directly and these probabilities allow us to expand the translation model for the new words.

There are other related works (Rapp, 1999; Daumé III and Jagarlamudi, 2011; Durrani and Koehn, 2014) that have explored approaches based on extracting lexicons using corpus based methods to resolve out of training vocabulary problems. There is also a rich body of recent literature that focuses on obtaining bilingual word embeddings using either sentence aligned corpora or document aligned corpora (Bhattacharai, 2012; Gouws et al., 2015; Kočiský et al., 2014). Our approach is significantly different as we obtain embeddings separately on monolingual corpora and then use supervision in the form of a small sparse bilingual dictionary.

## 3 Mapping Continuous Word Representations using a Bilinear Model

**Definitions.** Let  $\mathcal{E}$  and  $\mathcal{F}$  be the vocabularies of the two languages, source and target, and let  $e \in \mathcal{E}$  and  $f \in \mathcal{F}$  be the words in the languages respectively. We are given with a relatively small set of source word to target word  $e \rightarrow f$  dictionary. We also assume that we have access to some kind of distributed word embeddings in both languages. Let  $\phi(\cdot) \rightarrow \mathbb{R}^n$  denote the  $n$ -dimensional

distributed representation of the words, and let us assume we have both source ( $\phi_s$ ) and target ( $\phi_t$ ) embeddings. The task we are interested in is to learn a model for the conditional probability distribution  $\Pr(e|f)$ . That is, given a word in a source language, say English ( $e$ ), we want to get a conditional probability distribution of all the words in a foreign language ( $f$ ).

**Log-Bilinear Softmax Model.** We look at this task as a bilinear prediction task as proposed by Madhyastha et al. (2014b) and extend it for the bilingual setting. The proposed model makes use of word embeddings on both languages with no additional features. The basic function is formulated as log-bilinear softmax model and takes the following form:

$$\Pr(f|e; W) = \frac{\exp\{\phi_s(e)^\top W \phi_t(f)\}}{\sum_{f' \in \mathcal{F}} \exp\{\phi_s(e)^\top W \phi_t(f')\}} \quad (1)$$

Essentially, our problem reduces to: a) first getting the corresponding word embeddings of the vocabularies on both the languages on a significantly large monolingual corpus and b) estimating  $W$  given a relatively small dictionary. That is, to learn  $W$  we use the source word to target word dictionary as training supervision.

We learn  $W$  by minimizing the negative log-likelihood of the dictionary using a nuclear norm regularized objective as:  $L(W) = -\sum_{s,t} \log(\Pr(t|s; W)) + \lambda \|W\|_p$ .  $\lambda$  is the constant that controls the capacity of  $W$ . To find the optimum, we follow the previous work and use an optimization scheme based on Forward-Backward Splitting (FOBOS) (Singer and Duchi, 2009). We experiment with two regularization schemes,  $p = 2$  or the euclidean norm and  $p = *$  or the trace norm. In our experiments we found that both the norms have approximately similar performance, however the trace norm regularized  $W$  has lower capacity and hence, is less number of parameters. This is also observed by (Bach, 2008; Madhyastha et al., 2014b; Madhyastha et al., 2014a).

A by-product of regularizing with trace norm is that we obtain low-dimensional, aligned-compressed embeddings for both languages. This is possible because of the induced low-dimensional properties of  $W$ . That is, assume  $W$  has rank  $k$ , where  $k < n$ , such that  $W \approx U_k V_k^\top$ , then the product  $\phi_s(e)^\top U_k V_k^\top \phi_t(f)$  gives

us  $\phi_s(e)^\top U_k$  and  $V_k^\top \phi_t(f)$  compressed embeddings with shared properties. We leave exploration of the compressed embeddings for future work.

## 4 Experiments

**Data and System Settings.** For estimating the word embeddings we use the CBOV algorithm as implemented in the `Word2Vec` package (Mikolov et al., 2013a)<sup>1</sup> using a 5-token window. We obtain 300 dimension vectors for English and Spanish from a Wikipedia dump of 2015<sup>2</sup>, and the Quest data<sup>3</sup> which includes subcorpora such as United Nations and Europarl. The final corpus contains 2.27 billion tokens for English and 840 million tokens for Spanish. We obtain a coverage of 97% of the words in our test sets. We also remove any occurrence of sentences from the test set that are contained in our corpus, and avoid any transduction based knowledge transfer.

To train the log-bilinear softmax based model, we use the dictionary from the `Apertium project`<sup>4</sup> (Forcada et al., 2011). The dictionary contains 37651 words, we used 70% of them for training the log-bilinear model and 30% as a development set for model selection. The average precision @1 was 85.66% for the best model over the dev set.

On the other hand, we build a state-of-the-art phrase-based SMT system trained on the standard Europarl corpus for the English-to-Spanish language pair. We use a 5-gram language model that is estimated on the target side of the corpus using interpolated Kneser-Ney discounting with `SRILM` (Stolcke, 2002). Additional monolingual data available within Quest corpora is used to build a larger language model with the same characteristics. Word alignment is done with `GIZA++` (Och and Ney, 2003) and both phrase extraction and decoding are done with the `Moses` package (Koehn et al., 2007).

At decoding time, `Moses` allows to include additional translation pairs with their associated probabilities to selected words via xml mark-up. We take advantage of this feature to add our probabilistic estimations to each OOV. Since, by definition, OOV words do not appear in the parallel

<sup>1</sup><https://code.google.com/archive/p/word2vec/>

<sup>2</sup>Dumps downloaded in January 2015 from <https://dumps.wikimedia.org>.

<sup>3</sup><http://goo.gl/72LLXN>

<sup>4</sup>The bilingual dictionary can be downloaded here: <http://goo.gl/TjH31q>.

Table 1: OOVs on the dev and test sets.

	Sent.	Tokens	OOV <sub>all</sub>	OOV <sub>CW</sub>
NewsDev	3003	72988	1920 (2.6%)	378 (0.5%)
NewsTest	3000	64810	1590 (2.5%)	296 (0.5%)
WikiTest	500	11069	798 (7.2%)	201 (1.8%)

training corpus, they are not present in the translation model either and the new translation options only interact with the language model.

The optimization of the weights of the model with the additional translation options is trained with MERT (Och, 2003) against the BLEU (Papineni et al., 2002) evaluation metric on the NewsCommentaries 2012<sup>5</sup> (NewsDev) set. We test our systems on the NewsCommentaries 2013 set (NewsTest) for an in-domain evaluation and on a test set extracted from Wikipedia by Smith *et al.* (2010) for an out-of-domain evaluation (WikiTest).

The *domainness* of the test set is established with respect to the number of OOVs. Table 1 shows the figures of these sets paying special attention to the OOVs in the basic SMT system. Less than a 3% of the tokens are OOVs for News data (OOV<sub>all</sub>), whereas it is more than a 7% for Wikipedia’s. In our experiments, we distinguish between OOVs that are named entities and the rest of content words (OOV<sub>CW</sub>). Only about 0.5% (NewsTest) and 1.8% (WikiTest) of the tokens fall into this category, but we show that they are relevant for the final performance.

**Evaluation.** We consider two baseline systems, the first one does not output any translation for OOVs (*noOOV*), it just ignores the token; the second one outputs a verbatim copy of the unseen word as a translation (*verbatimOOV*). Table 2 shows the performance of these systems under three widely used evaluation metrics TER (Snover et al., 2006), BLEU and METEOR-paraphrase (MTR) (Banerjee and Lavie, 2005). Including the verbatim copy improves all the lexical evaluation metrics. Specially for named entities and acronyms (the 80% of OOVs in our sets), this is a hard baseline to beat since in most cases the same word is the correct translation (e.g. Messi, PHP, Sputnik...).

Next, we enrich the systems with information gathered from the large monolingual corpora in two ways, using a bigger language model (*BLM*)

<sup>5</sup><http://www.statmt.org/wmt13/translation-task.html>

Table 2: Automatic evaluation of the translation systems defined in Section 2. The best system is bold-faced (see text for statistical significance).

	NewsTest			WikiTest		
	TER	BLEU	MTR	TER	BLEU	MTR
noOOV	58.21	21.94	45.79	61.26	16.24	38.76
verbatimOOV	57.90	22.89	47.06	58.55	21.90	45.77
BWE	58.33	22.23	45.76	58.38	21.96	44.84
BWE <sub>CW50</sub>	57.66	23.09	47.14	56.19	24.16	48.49
BWE <sub>CW10</sub>	57.85	23.06	47.11	55.64	24.71	49.05
BLM	55.37	25.83	<b>49.19</b>	52.60	30.63	51.04
BLM+BWE	55.89	24.92	47.84	51.02	32.20	52.09
BLM+BWE <sub>50</sub>	55.55	25.61	49.01	49.50	33.94	54.93
BLM+BWE <sub>10</sub>	<b>55.31</b>	<b>25.86</b>	49.04	<b>49.12</b>	<b>34.58</b>	<b>55.52</b>

and using our newly proposed log-bilinear model that uses word embeddings (*BWE*). BLMs are very important to improve the fluency of the translations, however they may not be helpful for resolving out-of-vocabulary words. On the other hand, BWEs are important to make available to the decoder new vocabulary on the topic of the otherwise OOVs. Given the large percentage of named entities in the test sets (Table 1), our models add the source word as an additional option to the list of target words to mimic the *verbatimOOV* system.

Table 2 includes seven systems with the additional monolingual information. Three of them add, at decoding time, the top-*n* translation options given by the BWE for a OOV. *BWE* system uses the top-50 for all the OOVs, *BWE<sub>CW50</sub>* also uses the top-50 but only for content words other than named entities<sup>6</sup>, and *BWE<sub>CW10</sub>* limits the list to 10 elements. *BLM* is the same as the baseline system *verbatimOOV* but with the large language model. *BLM+BWE*, *BLM+BWE<sub>50</sub>* and *BLM+BWE<sub>10</sub>* combine the three BWE systems with the large language model.

A large number of unseen words in the NewsTest are mostly named entities, using BWEs to translate all the words, including named entities, barely improves the translation. Also, the richness in vocabulary, consisting of many names, adds noise to the decoder. We observe that the improvements are moderate in the NewsTest (in-domain dataset), this mostly is because the differences in the probability of the BWE translation options are very small owing to the candidates being named entities. We also see that this affects the overall integration of the scores into the decoder and

<sup>6</sup>We consider a named entity any word that begins with a capital letter and is not after a punctuation mark, and any fully capitalized word.

Table 3: Top-*n* list of translations obtained with the bilingual embeddings.

GALAXY	NYPHYS	STUART	FOLKSONG
<b>galaxia</b>	<b>ninfas</b>	William	música
planeta	ninfa	Henry	<b>folclore</b>
universo	crías	John	literatura
planetas	diosa	Charles	himno
galaxias	dioses	Thomas	folklore
...	...	Estuardo (#48)	<b>canción (#7)</b>

also induces ambiguity in the system. On the other hand, we observe that the decoder benefits from the information on content words, specially for the out-of-domain WikiTest set, given the constrained list of alternative translations (*BWE<sub>CW10</sub>* achieves 2.75 BLEU points of improvement).

The addition of the large language model improves the results significantly. When combined with the BWEs we observe that the BWEs clearly help in the translation of WikiTest but do not seem as relevant in the in-domain set. We also achieve a statistically significant improvement of 3.9 points of BLEU with the BLM and BWE combo system in WikiTest ( $p < 0.001$ ). The number of translation options in the list is also relevant, we see that for *BWE<sub>CW50</sub>* we have an improvement of 3.3 points on BLEU. We also observe that the results are consistent among different metrics.

We have further manually evaluated the translation of WikiTest using *BWE<sub>CW50</sub>*. We obtained an accuracy of a 68%, that is, the BWE gives the correct translation option at least 68% of the times. The other 32% of time, it fails as the words in the translated language happened to be either multiwords or named entities. In table 3 we observe some of the these examples. The first two examples *galaxy* and *nymphs* are nouns where we obtain the first option as the correct translation. The problem is harder for named entities as we observe in the table, the name *Stuart* in English has *William* as most probable translation in Spanish, the correct translation *Estuardo* however appears as the 48th choice. Our model is also unable to generate multiword expressions, as shown in the table for the english word *folksong*, the correct translation being *canción folk*. This would need two words in Spanish in order to be translated properly, however, our model does obtain words: *canción* and *folclore* as the most probable translation options.

## 5 Conclusions

We have presented a method for resolving unseen words in SMT that performs vocabulary expansion by using a simple log-bilinear softmax based model. The addition of translation options to a mere 1.8% of the words has allowed the system to obtain a relative improvement of a 13% in BLEU (3.9 points) for out-of-domain data. For in-domain data, where the number of content words is small, improvements are moderate. We would like to further study the repercussion of this simple method on diverse and most distant language pairs and how the form of the loss function affects the quality of the bilingual word embeddings.

## References

- [Al-Onaizan and Knight2002] Yaser Al-Onaizan and Kevin Knight. 2002. Machine transliteration of names in arabic text. In *Proceedings of the ACL-02 workshop on Computational approaches to semitic languages*, pages 1–13. Association for Computational Linguistics.
- [Bach2008] Francis R Bach. 2008. Consistency of the group lasso and multiple kernel learning. *The Journal of Machine Learning Research*, 9:1179–1225.
- [Banerjee and Lavie2005] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan, June. Association for Computational Linguistics.
- [Bhatarai2012] Alexandre Klementiev Ivan Titov Binod Bhatarai. 2012. Inducing Crosslingual Distributed Representations of Words. In *Proceedings of COLING 2012*, pages 1459–1474, Mumbai, India, December. The COLING 2012 Organizing Committee.
- [Callison-Burch et al.] Chris Callison-Burch, Philipp Koehn, and Miles Osborne. Improved Statistical Machine Translation Using Paraphrases. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*.
- [Daumé III and Jagarlamudi2011] Hal Daumé III and Jagadeesh Jagarlamudi. 2011. Domain adaptation for machine translation by mining unseen words. In *Association for Computational Linguistics*, Portland, OR.
- [Dou and Knight2012] Qing Dou and Kevin Knight. 2012. Large scale decipherment for out-of-domain machine translation. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 266–275. Association for Computational Linguistics.
- [Durrani and Koehn2014] Nadir Durrani and Philipp Koehn. 2014. Improving machine translation via triangulation and transliteration. In *Proceedings of the 17th Annual Conference of the European Association for Machine Translation (EAMT)*, Dubrovnik, Croatia.
- [Forcada et al.2011] Mikel L Forcada, Mireia Ginestí-Rosell, Jacob Nordfalk, Jim O’Regan, Sergio Ortiz-Rojas, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Gema Ramírez-Sánchez, and Francis M Tyers. 2011. Apertium: a free/open-source platform for rule-based machine translation. *Machine translation*, 25(2):127–144.
- [Gouws et al.2015] Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. BilBOWA: Fast Bilingual Distributed Representations without Word Alignments. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 748–756, July.
- [Habash2008] Nizar Habash. 2008. Four Techniques for Online Handling of Out-of-Vocabulary Words in Arabic-English Statistical Machine Translation. In *ACL 2008, Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics*, pages 57–60, June.
- [Hermjakob et al.2008] Ulf Hermjakob, Kevin Knight, and Hal Daumé III. 2008. Name translation in statistical machine translation-learning when to transliterate.
- [Irvine and Callison-Burch2013] Ann Irvine and Chris Callison-Burch. 2013. Combining bilingual and comparable corpora for low resource machine translation. In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pages 262–270.
- [Ishiwatari et al.2016] Shonosuke Ishiwatari, Naoki Yoshinaga, Masashi Toyoda, and Masaru Kitsuregawa. 2016. Instant translation model adaptation by translating unseen words in continuous vector space.
- [Kočíský et al.2014] Tomáš Kočíský, Karl Moritz Hermann, and Phil Blunsom. 2014. Learning bilingual word representations by marginalizing alignments. *arXiv preprint arXiv:1405.0947*.
- [Koehn et al.2007] Philipp Koehn, Hieu Hoang, Alexandra Birch Mayne, Christopher Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Annual Meeting of the Association for Computational Linguistics (ACL), Demonstration Session*, pages 177–180, June.

- [Madhyastha et al.2014a] Pranava Swaroop Madhyastha, Xavier Carreras, and Ariadna Quattoni. 2014a. Tailoring word embeddings for bilexical predictions: An experimental comparison. *International Conference on Learning Representations 2015, Workshop Track*.
- [Madhyastha et al.2014b] Swaroop Pranava Madhyastha, Xavier Carreras, and Ariadna Quattoni. 2014b. Learning task-specific bilexical embeddings. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 161–171. Dublin City University and Association for Computational Linguistics.
- [Marton et al.2009] Yuval Marton, Chris Callison-Burch, and Philip Resnik. 2009. Improved statistical machine translation using monolingually-derived paraphrases. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, pages 381–390. Association for Computational Linguistics.
- [Mikolov et al.2013a] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient Estimation of Word Representations in Vector Space. In *Proceedings of Workshop at ICLR*. <http://code.google.com/p/word2vec>.
- [Mikolov et al.2013b] Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013b. Exploiting Similarities among Languages for Machine Translation. [abs/1309.4168](https://arxiv.org/abs/1309.4168).
- [Och and Ney2003] Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.
- [Och2003] Franz Josef Och. 2003. Minimum Error Rate Training in Statistical Machine Translation. In *Proceedings of the Association for Computational Linguistics*, pages 160–167, Sapporo, Japan, July 6–7.
- [Papineni et al.2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the Association of Computational Linguistics*, pages 311–318.
- [Rapp1999] Reinhard Rapp. 1999. Automatic identification of word translations from unrelated english and german corpora. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, pages 519–526. Association for Computational Linguistics.
- [Saluja et al.2014] Avneesh Saluja, Hany Hassan, Kristina Toutanova, and Chris Quirk. 2014. Graph-based Semi-Supervised Learning of Translation Models from Monolingual Data. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014*, pages 676–686, June.
- [Singer and Duchi2009] Yoram Singer and John C Duchi. 2009. Efficient learning using forward-backward splitting. In *Advances in Neural Information Processing Systems*, pages 495–503.
- [Smith et al.2010] Jason R. Smith, Chris Quirk, and Kristina Toutanova. 2010. Extracting Parallel Sentences from Comparable Corpora Using Document Level Alignment. In *In Proceedings of Human Language Technologies: The 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT 2010)*, pages 403–411.
- [Snover et al.2006] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul. 2006. A Study of Translation Edit Rate with Targeted Human Annotation. In *Proceedings of the Seventh Conference of the Association for Machine Translation in the Americas (AMTA 2006)*, pages 223–231, Cambridge, Massachusetts, USA.
- [Stolcke2002] Andreas Stolcke. 2002. SRILM - An Extensible Language Modeling Toolkit. In *Proceedings of the Seventh International Conference of Spoken Language Processing (ICSLP2002)*, pages 901–904, Denver, Colorado, USA.
- [Wu et al.2014] Haiyang Wu, Daxiang Dong, Xiaoguang Hu, Dianhai Yu, Wei He, Hua Wu, Haifeng Wang, and Ting Liu. 2014. Improve Statistical Machine Translation with Context-Sensitive Bilingual Semantic Embedding Model. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 142–146, Doha, Qatar, October. Association for Computational Linguistics.
- [Zhao et al.2015] Kai Zhao, Hany Hassan, and Michael Auli. 2015. Learning Translation Models from Monolingual Continuous Representations. In *The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL HLT 2015*, pages 1527–1536, June.
- [Zou et al.2013] Will Y. Zou, Richard Socher, Daniel M. Cer, and Christopher D. Manning. 2013. Bilingual Word Embeddings for Phrase-Based Machine Translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013*, pages 1393–1398, October.