

Machine Translation of Low-Resource Indo-European Languages

Wei-Rui Chen Muhammad Abdul-Mageed

Natural Language Processing Lab
The University of British Columbia

{weirui.chen,muhammad.mageed}@ubc.ca

Abstract

In this work, we investigate methods for the challenging task of translating between low-resource language pairs that exhibit some level of similarity. In particular, we consider the utility of transfer learning for translating between several Indo-European low-resource languages from the Germanic and Romance language families. In particular, we build two main classes of transfer-based systems to study how relatedness can benefit the translation performance. The *primary* system fine-tunes a model pre-trained on a *related* language pair and the *contrastive* system fine-tunes one pre-trained on an *unrelated* language pair. Our experiments show that although relatedness is not necessary for transfer learning to work, it does benefit model performance.

1 Introduction

Machine translation (MT) is currently one of the hot application areas of deep learning, with neural machine translation (NMT) achieving outstanding performance where large amounts of parallel data are available (Koehn and Knowles, 2017; Ranathunga et al., 2021; Luong et al., 2015a; Nguyen and Chiang, 2017). In low-resource settings, transfer learning methods have proven useful for improving system performance (Zoph et al., 2016; Nguyen and Chiang, 2017; Kocmi and Bojar, 2018). In this work, we focus on studying NMT, including in the low-resource scenario. In particular, we focus our attention on investigating the effect of *language relatedness* on the transfer process. We define relatedness of a pair of languages based on belonging to the same language family. That is, by ‘related’ we mean ‘within the *same* language family’ whereas by ‘unrelated’ we mean ‘belong to two *different* language families’. For example, we call English and Swedish related since they belong to the Germanic language family but English and French *not* related since the latter

belongs to the Romance language family. As an analogy to human learning, we would like to ask: *if there are two translators (pre-trained models), for example one Catalan → Spanish translator and one Catalan → English translator, will they (after extra training, i.e., fine-tuning/transfer learning process) have different abilities to translate from Catalan into Occitan?* If the Catalan-Spanish translator proves to perform better Catalan → Occitan, we may attribute this to Spanish and Occitan being members of the Romance language family while English being a member of the, Germanic, different family.

Of particular interest to us are two sets of languages belonging to two different language families, one set to Romance and the other set to Germanic. For the former set, we take Catalan (ca), Italian (it), Occitan (oc), Romanian (ro), and Spanish (es); and we take English (en) for the latter set. We note that both Romance and Germanic are two branches of the larger Indo-European language family, and hence there are some level of relatedness between all the languages we study in this work. Nevertheless, languages in Romance and Germanic differ in some syntactic structures. For example, the position of attributive adjectives in Germanic languages is before the noun while it is after the noun for Romance languages (Van de Velde et al., 2014). Despite differences, the writing system of all languages in this work is the Latin script. This can be beneficial to transfer learning because these languages can potentially share common lexical items or morphemes, which may facilitate the transfer learning process.

As mentioned, we adopt transfer learning since it has been shown to improve translation quality for low-resource language pairs. For example, Zoph et al. (2016) sequentially build a *parent* model and a *child* model where each is trained, respectively, on high-resource and low-resource language pairs with the child model retaining parameters from

Lang. Pair	Primary	Contrastive
ca-it	ca-it/ca-es	ca-en
ca-oc	ca-es	ca-en
ca-ro	ca-es	ca-en

Table 1: Pre-trained model choices for our primary and contrastive NMT systems.

the parent model. In addition, [Nguyen and Chiang \(2017\)](#) successfully transfer knowledge from a parent model to a child model where both models are trained on low-resource but related language pairs. [Kocmi and Bojar \(2018\)](#) also adopt a similar approach to these previous works, but based their work on the Transformer architecture ([Vaswani et al., 2017](#)) instead of using a recurrent encoder-decoder network with attention ([Bahdanau et al., 2015](#); [Luong et al., 2015b](#)).

Our work builds on these studies. Namely, we empirically apply transfer learning under different conditions to members of various language families. The outcomes of our work are similar to those of [Zoph et al. \(2016\)](#); [Kocmi and Bojar \(2018\)](#). That is, while we find relatedness to be beneficial, a positive transfer between an unrelated language pair can still be possible (although with a potentially diminished performance).

The rest of this paper is organized as follows: In Section 2, we overview related work. We introduce our datasets and experiments in Section 3. In Section 4, we present and analyze our results. We conclude in Section 5.

2 Background

2.1 Transfer learning

Transfer learning is a machine learning approach that aims at transferring the knowledge of one task to another. As an analogy to human learning, one who masters the skills to ride a bicycle may transfer the knowledge to riding a motorcycle because these two tasks share common abilities such as maintaining balance on a two-wheel moving vehicle ([Pan and Yang, 2009](#); [Weiss et al., 2016](#)). We employ transfer learning to port knowledge from a model trained on one pair of languages to another. We now discuss transfer learning in NMT.

2.2 Transfer learning in Machine Translation

[Zoph et al. \(2016\)](#) design a framework where a parent model is trained on a high-resource language pair while retaining model parameters for

the child model to start fine-tuning with. Using this method, [Zoph et al. \(2016\)](#) improve system performance by an average of 5.6 BLEU points. The improvement is realized by transferring what is learnt in the high-resource language pair to the low-resource language pair. The Uzbek-English model obtains 10.7 BLEU score without the parent model and improves to 15.0 with the French-English parent model. The Spanish-English model has 16.4 BLEU score without the parent model and 31.0 with the French-English parent model. These results show that applying transfer learning contributes 4.3 and 14.6 BLEU points gain. Based on results from [Zoph et al. \(2016\)](#), the closer the two source languages, the more performance gain acquired. Due to the relatedness between Spanish and French (both are members of the Roman language family), performance gain is higher for this pair.

Following previous work, [Nguyen and Chiang \(2017\)](#) design a paradigm similar to that of [Zoph et al. \(2016\)](#) but maintain one major difference. In particular, [Nguyen and Chiang \(2017\)](#) try to make use of relatedness between the parent and child models at the vocabulary level: instead of randomly mapping tokens in the parent and child vocabulary, they retain the parent tokens for the child model if these tokens exist in child language pair. This approach is based on two assumptions - (i) the lexicons of the parent and child language pair have at least some partial overlap and (ii) these identical tokens have similar meaning. Instead of the word-level tokenization in [Zoph et al. \(2016\)](#), [Nguyen and Chiang \(2017\)](#) use Byte Pair Encoding (BPE) ([Gage, 1994](#); [Sennrich et al., 2016](#)) to obtain subword tokens which may increase the number of overlapped tokens between the parent and child models. Improvement of 0.8 and 4.3 in BLEU score were obtained for the Turkish-English and Uyghur-English child models as transferred from an Uzbek-English parent model.

Following the previous two works, [Kocmi and Bojar \(2018\)](#) take a similar approach but use the Transformer architecture. They obtain an improvement of 3.38 BLEU for an English-Estonian child model transferred from an English-Czech parent model. Similarly, [Neubig and Hu \(2018\)](#) add a second language related to the added low-resource language to avoid overfitting when fine-tuning. This mechanism has shown to be effective. Other works have investigated NMT approaches

to similar languages by pre-training new language models on the low-resource languages (Nagoudi et al., 2021) or without necessarily applying transfer learning (Przystupa and Abdul-Mageed, 2019; Adebara et al., 2020; Barrault et al., 2019, 2020), and there are several works on low resource languages (Adebara et al., 2021). We now introduce our experimental settings.

3 Experimental Settings

3.1 Languages & Settings

We carry out experiments on three language pairs: *ca-it*, *ca-oc*, and *ca-ro*. The number of parallel sentences of each dataset is shown in Table 4. Training data are from OPUS (Tiedemann, 2012), particularly version 1 of the WikiMatrix datasets (Schwenk et al., 2021). They are data of child language pair and are used to fine-tune pre-trained model. Development and test data are provided by organizers of the Multilingual Low-Resource Translation for Indo-European Languages shared task. The shared task is hosted in EMNLP 2021 Sixth Conference on Machine Translation (WMT21).

We build two systems: a *Primary system* and a *Contrastive system*. The primary system fine-tunes pre-trained *ca-es* and *ca-it* models, while the contrastive system fine-tunes a pre-trained *ca-en* model as shown in Table 1. The primary and contrastive systems serve as context for studying the role of language *relatedness* in transfer learning in NMT. We submitted predictions of the two systems to the WMT2021 shared task, and evaluation was based on blind test sets with a number of metrics as run by shared task organizers. An exception is the *ca-it* language pair fine-tuned on top of the *ca-es* pre-trained model—we train the model of this pair post shared task formal evaluation.

3.2 Model Architecture

We leverage publicly accessible pre-trained models on Huggingface (Wolf et al., 2020) from Helsinki-NLP (Tiedemann and Thottingal, 2020). The pre-trained MT models released by Helsinki-NLP are trained on OPUS. These models are Transformer-based implemented in the Marian-NMT framework (Junczys-Dowmunt et al., 2018). Each model has six self-attention layers in both the encoder and decoder sides, and each layer has eight attention heads. The tokenization method is SentencePiece (Kudo and Richardson, 2018) which

produces vocabulary of size 49,621, 21,528 and 55,255 for *ca-es*, *ca-it*, and *ca-en* models, respectively.

3.3 Approach

The pre-trained models are chosen based on the degree of *relatedness* of the original target language on which the model is trained and the new target language on which the model is fine-tuned. *Primary system* takes related languages while *contrastive system* takes unrelated languages. Since Catalan, Italian, Occitan, Romanian, and Spanish are all members of the Roman language family, we take **ca-es** as our pre-trained MT model for transfer learning. As English is a member of the Germanic language family, we use a **ca-en** pre-trained model for our transfer learning. Our model choices are summarized in Table 1.

Without modifying the architecture of the MT pre-trained models, all architecture-related hyperparameters are identical to the original edition. As for hyperparameters related to fine-tuning, the number of beams for beam search is modified from four for pre-training to six for fine-tuning. The batch size is set to be 25. Pre-trained models are further fine-tuned for 30,000 steps on OPUS bitext. The checkpoint with the lowest validation loss is then selected as our best model for prediction.

Similar to Zoph et al. (2016); Nguyen and Chiang (2017); Kocmi and Bojar (2018), to achieve transfer learning, we retain the parameters of the parent model when fine-tuning the child model. Besides, parent and child models share a common vocabulary. That is, we do not build distinct vocabularies for the parent model and child models. A shared vocabulary can contribute to better transfer learning since all our language pairs employ the same Latin writing system. We suspect a shared vocabulary is more influential when the two languages are related to each other since the languages may have common morphemes, lexical items, or syntactical structure. For unrelated languages, a shared vocabulary may not hurt since the token embeddings are not frozen throughout the fine-tuning process. That is, token embeddings can still be updated to attain better representations during training.

3.4 Baseline Models

To demonstrate the effectiveness of our transfer learning approach, we provide a baseline model for each language pair that is simply a parent model (a

		Baseline		Our models				
Pre-trained Model	Lang. Pairs	BLEU	chrF	BLEU	chrF	TER	COMET	BertScore
ca-it	ca-it	29.31	0.583	35.06	0.622	0.477	0.391	0.886
ca-es	ca-it	7.07	0.370	33.13	0.602	0.499	-	-
ca-es	ca-oc	12.56	0.472	59.93	0.787	0.254	0.538	0.928
ca-es	ca-ro	4.43	0.266	11.24	0.354	0.855	-0.908	0.749

Table 2: Primary system results. We did not submit the ca-it language pair fine-tuned on the ca-es pre-trained model to the WMT2021 shared task, and hence the results are calculated by ourselves with Sacrebleu.

		Baseline		Our models				
Pre-trained Model	Lang. Pairs	BLEU	chrF	BLEU	chrF	TER	COMET	BertScore
ca-en	ca-it	1.97	0.249	25.46	0.539	0.574	-0.263	0.844
ca-en	ca-oc	2.16	0.258	51.46	0.736	0.316	0.259	0.905
ca-en	ca-ro	1.59	0.209	8.61	0.311	0.884	-1.119	0.725

Table 3: Contrastive system results

Lang. Pairs	Train	Dev	Test
ca-it	1,143,531	1269	1743
ca-oc	138,743	1269	1743
ca-ro	490,064	1269	1743

Table 4: Distribution of dataset

pre-trained model) without any fine-tuning on data of the child language pair.

3.5 Evaluation

The adopted metrics are BLEU (Papineni et al., 2002), chrF (Popović, 2015), TER (Olive, 2005), COMET (Rei et al., 2020), and BERTScore (Zhang et al., 2019). BLEU, chrF and TER are measured with the implementation of *Sacrebleu* (Post, 2018)¹.

4 Results and Analysis

4.1 Primary and Contrastive Systems

As can be seen in the rightmost five columns in Table 2 and Table 3, primary system outperforms contrastive system across all metrics. We believe that **ca-es** pre-trained MT model performs better transfer learning because Spanish is closer to Italian, Occitan, and Romanian than English is to these languages. These results, as such, indicate that transfer learning between related language pairs

can produce better performance than between unrelated language pairs.

4.2 Baseline and Fine-tuned Models

Our results in Table 2 and Table 3 show the effectiveness of transfer learning for both related and *unrelated* language pairs. This is the case since both systems experience a performance gain after fine-tuning.

As an interesting observation, it seems counter-intuitive to have the unrelated language pairs experience slightly higher performance gain. For example, regarding **ca-oc** language pair, the transfer learning provides 47.37 BLEU score improvement transferring from **ca-es** parent model but 49.3 BLEU score improvement transferring from **ca-en** parent model. We suspect this is because in our work, when fine-tuning, we fix source language and alter the target language.

Unlike multilingual MT models which requires target language label to be prepended at the beginning of a source sentence (Johnson et al., 2017) or notifying the model what target language is for this forward propagation (Liu et al., 2020), the pre-trained models we use in this work are bilingual models which lack a mechanism to provide the model any information about current target language. Therefore, the **ca-en** pre-trained model does not know it should now be translating Catalan to Occitan instead of English. Due to producing pre-

¹<https://github.com/mjpost/sacrebleu>

diction in an incorrect target language, the metrics will be very poor. After fine-tuning the parent models on data of the child language pairs, the models are likely able to produce prediction in the correct target language. Due to baseline metrics being too low, the difference in metric values between non-fine-tuned (baseline) and fine-tuned models are large and that is why the performance gain can be higher in contrastive system than in primary system.

5 Conclusion

In this work, we confirm previous works showing that transfer learning benefits NMT. Besides, an empirical comparison between transferring from related and unrelated languages shows that relatedness is not strictly required for knowledge transfer, but it does result in higher performance than transferring with unrelated languages.

Acknowledgements

We appreciate the support from the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada, Canadian Foundation for Innovation, Compute Canada (www.computecanada.ca), and UBC ARC-Sockeye (<https://doi.org/10.14288/SOCKEYE>).

References

- Ife Adebara, Muhammad Abdul-Mageed, and Miikka Silfverberg. 2021. [Translating the unseen? yoruba-english mt in low-resource, morphologically-unmarked settings](#). *arXiv preprint arXiv:2103.04225*.
- Ife Adebara, El Moatez Billah Nagoudi, and Muhammad Abdul Mageed. 2020. [Translating similar languages: Role of mutual intelligibility in multilingual transformers](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 381–386, Online. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015*.
- Loïc Barrault, Magdalena Biesialska, Ondrej Bojar, M. Costa-jussà, C. Federmann, Yvette Graham, Roman Grundkiewicz, B. Haddow, M. Huck, E. Jounis, Tom Kocmi, Philipp Koehn, Chi kiu Lo, Nikola Ljubesic, Christof Monz, Makoto Morishita, M. Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(wmt20\)](#). In *WMT@EMNLP*.
- Loïc Barrault, Ondřej Bojar, Marta R Costa-Jussa, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, et al. 2019. [Findings of the 2019 conference on machine translation \(wmt19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61.
- Philip Gage. 1994. [A new algorithm for data compression](#). *C Users Journal*, 12(2):23–38.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Tom Kocmi and Ondrej Bojar. 2018. [Trivial transfer learning for low-resource neural machine translation](#). *CoRR*, abs/1809.00357.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015a. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.

- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015b. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- El Moatez Billah Nagoudi, Wei-Rui Chen, Muhammad Abdul-Mageed, and Hasan Cavusogl. 2021. [Indt5: A text-to-text transformer for 10 indigenous languages](#). In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 265–271.
- Graham Neubig and Junjie Hu. 2018. [Rapid adaptation of neural machine translation to new languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 875–880, Brussels, Belgium. Association for Computational Linguistics.
- Toan Q. Nguyen and David Chiang. 2017. [Transfer learning across low-resource, related languages for neural machine translation](#). *CoRR*, abs/1708.09803.
- Joseph Olive. 2005. Global autonomous language exploitation (gale). *DARPA/IPTO Proposer Information Pamphlet*.
- Sinno Jialin Pan and Qiang Yang. 2009. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Michael Przystupa and Muhammad Abdul-Mageed. 2019. [Neural machine translation of low-resource and similar languages with backtranslation](#). In *Proceedings of the 4th Conference on MT (Volume 3: Shared Task Papers, Day 2)*, pages 224–235.
- Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2021. [Neural machine translation for low-resource languages: A survey](#).
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [Wiki-Matrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT — Building open translation services for the World](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.
- Freek Van de Velde, Petra Sleeman, and Harry Perriodon. 2014. The adjective in germanic and romance. *Adjectives in Germanic and Romance*, 212:1.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. 2016. A survey of transfer learning. *Journal of Big data*, 3(1):1–40.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.