

UDS–DFKI Submission to the WMT2019 Similar Language Translation Shared Task

Santanu Pal^{1,3}, Marcos Zampieri², Josef van Genabith^{1,3}

¹Department of Language Science and Technology, Saarland University, Germany

²Research Institute for Information and Language Processing, University of Wolverhampton, UK

³German Research Center for Artificial Intelligence (DFKI),
Saarland Informatics Campus, Germany
santanu.pal@uni-saarland.de

Abstract

In this paper we present the UDS-DFKI system submitted to the Similar Language Translation shared task at WMT 2019. The first edition of this shared task featured data from three pairs of similar languages: Czech and Polish, Hindi and Nepali, and Portuguese and Spanish. Participants could choose to participate in any of these three tracks and submit system outputs in any translation direction. We report the results obtained by our system in translating from Czech to Polish and comment on the impact of out-of-domain test data in the performance of our system. UDS-DFKI achieved competitive performance ranking second among ten teams in Czech to Polish translation.

1 Introduction

The shared tasks organized annually at WMT provide important benchmarks used in the MT community. Most of these shared tasks include English data, which contributes to make English the most resource-rich language in MT and NLP. In the most popular WMT shared task for example, the News task, MT systems have been trained to translate texts from and to English (Bojar et al., 2016, 2017).

This year, we have observed a shift on the dominant role that English on the WMT shared tasks. The News task featured for the first time two language pairs which did not include English: German-Czech and French-German. In addition to that, the Similar Language Translation was organized for the first time at WMT 2019 with the purpose of evaluating the performance of MT systems on three pairs of similar languages from three different language families: Ibero-Romance, Indo-Aryan, and Slavic.

The Similar Language Translation (Barrault et al., 2019) task provided participants with train-

ing, development, and testing data from the following language pairs: Spanish - Portuguese (Romance languages), Czech - Polish (Slavic languages), and Hindi - Nepali (Indo-Aryan languages). Participant could submit system outputs to any of the three language pairs in any direction. The shared task attracted a good number of participants and the performance of all entries was evaluated using popular MT automatic evaluation metrics, namely BLEU (Papineni et al., 2002) and TER (Snover et al., 2006).

In this paper we describe the UDS-DFKI system to the WMT 2019 Similar Language Translation task. The system achieved competitive performance and ranked second among ten entries in Czech to Polish translation in terms of BLEU score.

2 Related Work

With the widespread use of MT technology and the commercial and academic success of NMT, there has been more interest in training systems to translate between languages other than English (Costa-jussà, 2017). One reason for this is the growing need of direct translation between pairs of similar languages, and to a lesser extent language varieties, without the use of English as a pivot language. The main challenge is to overcome the limitation of available parallel data taking advantage of the similarity between languages. Studies have been published on translating between similar languages (e.g. Catalan - Spanish (Costa-jussà, 2017)) and language varieties such as European and Brazilian Portuguese (Fancellu et al., 2014; Costa-jussà et al., 2018). The study by Lakew et al. (2018) tackles both training MT systems to translate between European–Brazilian Portuguese and European–Canadian French, and two pairs of similar languages Croatian–Serbian

and Indonesian–Malay.

Processing similar languages and language varieties has attracted attention not only in the MT community but in NLP in general. This is evidenced by a number of research papers published in the last few years and the recent iterations of the VarDial evaluation campaign which featured multiple shared tasks on topics such as dialect detection, morphosyntactic tagging, cross-lingual parsing, cross-lingual morphological analysis (Zampieri et al., 2018, 2019).

3 Data

We used the Czech–Polish dataset provided by the WMT 2019 Similar Language Translation task organizers for our experiments. The released parallel dataset consists of out-of-domain (or general-domain) data only and it differs substantially from the released development set which is part of a TED corpus. The parallel data includes Europarl v9, Wiki-titles v1, and JRC-Acquis. We combine all the released data and prepare a large out-domain dataset.

3.1 Pre-processing

The out-domain data is noisy for our purposes, so we apply methods for cleaning. We performed the following two steps: (i) we use the cleaning process described in Pal et al. (2015), and (ii) we execute the Moses (Koehn et al., 2007) corpus cleaning scripts with minimum and maximum number of tokens set to 1 and 100, respectively. After cleaning, we perform punctuation normalization, and then we use the Moses tokenizer to tokenize the out-domain corpus with ‘no-escape’ option. Finally, we apply true-casing.

The cleaned version of the released data, i.e., the General corpus containing 1,394,319 sentences, is sorted based on the score in Equation 1. Thereafter, We split the entire data (1,394,319) into two sets; we use the first 1,000 for validation and the remaining as training data. The released development set (Dev) is used as test data for our experiment. It should be noted that, we exclude 1,000 sentences from the General corpus which are scored as top (i.e., more in-domain like) during the data selection process.

We prepare two parallel training sets from the aforementioned training data: (i) *transference500K* (presented next), collected 500,000 parallel data through data selection method (Axelrod

et al., 2011), which are very similar to the in-domain data (for our case the development set), and (ii) *transferenceALL*, utilizing all the released out-domain data sorted by Equation 1.

The *transference500K* training set is prepared using in-domain (development set) bilingual cross-entropy difference for data selection as was described in Axelrod et al. (2011). The difference in cross-entropy is computed based on two language models (LM): a domain-specific LM is estimated from the in-domain (containing 2050 sentences) corpus (lm_i) and the out-domain LM (lm_o) is estimated from the eScape corpus. We rank the eScape corpus by assigning a score to each of the individual sentences which is the sum of the three cross-entropy (H) differences. For a j^{th} sentence pair src_j – trg_j , the score is calculated based on Equation 1.

$$score = |H_{src}(src_j, lm_i) - H_{src}(src_j, lm_o)| + |H_{trg}(trg_j, lm_i) - H_{trg}(trg_j, lm_o)| \quad (1)$$

4 System Architecture - The Transference Model

Our *transference* model extends the original transformer model to multi-encoder based transformer architecture. The *transformer* architecture (Vaswani et al., 2017) is built solely upon such attention mechanisms completely replacing recurrence and convolutions. The transformer uses positional encoding to encode the input and output sequences, and computes both self- and cross-attention through so-called multi-head attentions, which are facilitated by parallelization. We use multi-head attention to jointly attend to information at different positions from different representation subspaces.

The first encoder (enc_1) of our model encodes word form information of the source (f_w), and a second sub-encoder (enc_2) to encode sub-word (byte-pair-encoding) information of the source (f_s). Additionally, a second encoder ($enc_{src \rightarrow mt}$) which takes the encoded representation from the enc_1 , combines this with the self-attention-based encoding of f_s (enc_2), and prepares a representation for the decoder (dec_e) via cross-attention. Our second encoder ($enc_{1 \rightarrow 2}$) can be viewed as a transformer based NMT’s decoding block, however, without masking. The intuition behind our

architecture is to generate better representations via both self- and cross-attention and to further facilitate the learning capacity of the feed-forward layer in the decoder block. In our transference model, one self-attended encoder for f_w , $\mathbf{f}_w = (w_1, w_2, \dots, w_k)$, returns a sequence of continuous representations, enc_1 , and a second self-attended sub-encoder for f_s , $\mathbf{f}_s = (s_1, s_2, \dots, s_l)$, returns another sequence of continuous representations, enc_2 . Self-attention at this point provides the advantage of aggregating information from all of the words, including f_w and f_s , and successively generates a new representation per word informed by the entire f_w and f_s context. The internal enc_2 representation performs cross-attention over enc_1 and prepares a final representation ($enc_{1 \rightarrow 2}$) for the decoder (dec_e). The decoder generates the e output in sequence, $\mathbf{e} = (e_1, e_2, \dots, e_n)$, one word at a time from left to right by attending to previously generated words as well as the final representations ($enc_{1 \rightarrow 2}$) generated by the encoder.

We use the scale-dot attention mechanism (like Vaswani et al. (2017)) for both self- and cross-attention, as defined in Equation 2, where Q , K and V are query, key and value, respectively, and d_k is the dimension of K .

$$attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (2)$$

The multi-head attention mechanism in the transformer network maps the Q , K , and V matrices by using different linear projections. Then h parallel heads are employed to focus on different parts in V . The i^{th} multi-head attention is denoted by $head_i$ in Equation 3. $head_i$ is linearly learned by three projection parameter matrices: $W_i^Q, W_i^K \in R^{d_{model} \times d_k}, W_i^V \in R^{d_{model} \times d_v}$; where $d_k = d_v = d_{model}/h$, and d_{model} is the number of hidden units of our network.

$$head_i = attention(QW_i^Q, KW_i^K, VW_i^V) \quad (3)$$

Finally, all the vectors produced by parallel heads are linearly projected using concatenation and form a single vector, called a multi-head attention (M_{att}) (cf. Equation 4). Here the dimension of the learned weight matrix W^O is $R^{d_{model} \times d_{model}}$.

$$M_{att}(Q, K, V) = Concat_{i:1}^n(head_i)W^O \quad (4)$$

5 Experiments

We explore our *transference* model –a two-encoder based transformer architecture, in CS-PL similar language translation task.

5.1 Experiment Setup

For *transferenceALL*, we initially train on the complete out-of-domain dataset (General). The General data is sorted based on their in-domain similarities as described in Equation 1.

transferenceALL models are then fine-tuned towards the 500K (in-domain-like) data. Finally, we perform checkpoint averaging using the 8 best checkpoints. We report the results on the provided development set, which we use as a test set before the submission. Additionally we also report the official test set result.

To handle out-of-vocabulary words and to reduce the vocabulary size, instead of considering words, we consider subword units (Sennrich et al., 2016) by using byte-pair encoding (BPE). In the preprocessing step, instead of learning an explicit mapping between BPEs in the Czech (CS) and Polish (PL), we define BPE tokens by jointly processing all parallel data. Thus, CS and PL derive a single BPE vocabulary. Since CS and PL belong to the similar language, they naturally share a good fraction of BPE tokens, which reduces the vocabulary size.

We pass word level information on the first encoder and the BPE information to the second one. On the decoder side of the transference model we pass only BPE text.

We evaluate our approach with development data which is used as test case before submission. We use BLEU (Papineni et al., 2002) and TER (Snober et al., 2006).

5.2 Hyper-parameter Setup

We follow a similar hyper-parameter setup for all reported systems. All encoders, and the decoder, are composed of a stack of $N_{fw} = N_{fs} = N_{es} = 6$ identical layers followed by layer normalization. Each layer again consists of two sub-layers and a residual connection (He et al., 2016) around each of the two sub-layers. We apply dropout (Srivastava et al., 2014) to the output of each sub-layer, before it is added to the sub-layer input and normalized. Furthermore, dropout is applied to the sums of the word embeddings and the corresponding positional encodings in both encoders as well

as the decoder stacks.

We set all dropout values in the network to 0.1. During training, we employ label smoothing with value $\epsilon_{ls} = 0.1$. The output dimension produced by all sub-layers and embedding layers is $d_{model} = 512$. Each encoder and decoder layer contains a fully connected feed-forward network (*FFN*) having dimensionality of $d_{model} = 512$ for the input and output and dimensionality of $d_{ff} = 2048$ for the inner layers. For the scaled dot-product attention, the input consists of queries and keys of dimension d_k , and values of dimension d_v . As multi-head attention parameters, we employ $h = 8$ for parallel attention layers, or heads. For each of these we use a dimensionality of $d_k = d_v = d_{model}/h = 64$. For optimization, we use the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 10^{-9}$.

The learning rate is varied throughout the training process, and increasing for the first training steps $warmup_{steps} = 8000$ and afterwards decreasing as described in (Vaswani et al., 2017). All remaining hyper-parameters are set analogously to those of the transformer’s *base* model. At training time, the batch size is set to 25K tokens, with a maximum sentence length of 256 subwords, and a vocabulary size of 28K. After each epoch, the training data is shuffled. After finishing training, we save the 5 best checkpoints which are written at each epoch. Finally, we use a single model obtained by averaging the last 5 checkpoints. During decoding, we perform beam search with a beam size of 4. We use shared embeddings between CS and PL in all our experiments.

6 Results

We present the results obtained by our system in Table 1.

tested on	model	BLEU	TER
Dev set	Generic	12.2	75.8
Dev set	Fine-tuned*	25.1	58.9
Test set	Generic	7.1	89.3
Test set	Fine-Tuned*	7.6	87.0

Table 1: Results for CS–PL Translation; * averaging 8 best checkpoints.

Our fine-tuned system on development set provides significant performance improvement over the generic model. We found +12.9 absolute BLEU points improvement over the generic

model. Similar improvement is also observed in terms of TER (-16.9 absolute). It is to be noted that our generic model is trained solely on the clean version of training data.

Before submission, we performed punctuation normalization, unicode normalization, and detokenization for the run.

In Table 2 we present the ranking of the competition provided by the shared task organizers. Ten entries were submitted by five teams and are ordered by BLEU score. TER is reported for all submissions which achieved BLEU score greater than 5.0. The type column specifies the type of system, whether it is a Primary (P) or Constrastive (C) entry.

Team	Type	BLEU	TER
UPC-TALP	P	7.9	85.9
UDS-DFKI	P	7.6	87.0
Uhelsinki	P	7.1	87.4
Uhelsinki	C	7.0	87.3
Incomslav	C	5.9	88.4
Uhelsinki	C	5.9	88.4
Incomslav	P	3.2	-
Incomslav	C	3.1	-
UBC-NLP	C	2.3	-
UBC-NLP	P	2.2	-

Table 2: Rank table for Czech to Polish Translation

Our system was ranked second in the competition only 0.3 BLEU points behind the winning team UPC-TALP. The relative low BLEU and high TER scores obtained by all teams are due to out-of-domain data provided in the competition which made the task equally challenging to all participants.

7 Conclusion

This paper presented the UDS-DFKI system submitted to the Similar Language Translation shared task at WMT 2019. We presented the results obtained by our system in translating from Czech to Polish. Our system achieved competitive performance ranking second among ten teams in the competition in terms of BLEU score. The fact that out-of-domain data was provided by the organizers resulted in a challenging but interesting scenario for all participants.

In future work, we would like to investigate how effective is the proposed hypothesis (i.e., word-BPE level information) in similar language trans-

lation. Furthermore, we would like to explore the similarity between these two languages (and the other two language pairs in the competition) in more detail by training models that can best capture morphological differences between them. During such competitions, this is not always possible due to time constraints.

Acknowledgments

This research was funded in part by the German research foundation (DFG) under grant number GE 2819/2-1 (project MMPE) and the German Federal Ministry of Education and Research (BMBF) under funding code 01IW17001 (project DeepLee). The responsibility for this publication lies with the authors. We would like to thank the anonymous WMT reviewers for their valuable input, and the organizers of the shared task.

References

- Amitai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. Domain Adaptation via Pseudo In-domain Data Selection. In *Proceedings of EMNLP*.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 Conference on Machine Translation (WMT19). In *Proceedings of WMT*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. Findings of the 2017 Conference on Machine Translation (WMT17). In *Proceedings of WMT*.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 Conference on Machine Translation. In *Proceedings of WMT*.
- Marta R. Costa-jussà. 2017. Why Catalan-Spanish Neural Machine Translation? Analysis, Comparison and Combination with Standard Rule and Phrase-based Technologies. In *Proceedings of VarDial*.
- Marta R. Costa-jussà, Marcos Zampieri, and Santanu Pal. 2018. A Neural Approach to Language Variety Translation. In *Proceedings of VarDial*.
- Federico Fancellu, Andy Way, and Morgan O'Brien. 2014. Standard Language Variety Conversion for Content Localisation via SMT. In *Proceedings of EAMT*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. *Proceedings of CVPR*.
- Diederik P Kingma and Jimmy Lei Ba. 2015. Adam: A Method for Stochastic Optimization. *Proceedings of ICLR*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open Source Toolkit for Statistical Machine Translation. In *Proceedings of ACL*.
- Surafel M Lakew, Aliia Erofeeva, and Marcello Federico. 2018. Neural Machine Translation into Language Varieties. *arXiv preprint arXiv:1811.01064*.
- Santanu Pal, Sudip Naskar, and Josef van Genabith. 2015. UdS-sant: English–German hybrid machine translation system. In *Proceedings of WMT*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of ACL*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of ACL*.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A Study of Translation Edit Rate with Targeted Human Annotation. In *Proceedings of AMTA*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958.
- 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 *Proceedings of NIPS*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Ahmed Ali, Suwon Shon, James Glass, Yves Scherrer, Tanja Samardžić, Nikola Ljubešić, Jörg Tiedemann, Chris van der Lee, Stefan Grondelaers, Nelleke Oostdijk, Dirk Speelman, Antal van den Bosch, Ritesh Kumar, Bornini Lahiri, and Mayank Jain. 2018. Language Identification and Morphosyntactic Tagging: The Second VarDial Evaluation Campaign. In *Proceedings of VarDial*.
- Marcos Zampieri, Shervin Malmasi, Yves Scherrer, Tanja Samardžić, Francis Tyers, Miikka Silfverberg, Natalia Klyueva, Tung-Le Pan, Chu-Ren Huang, Radu Tudor Ionescu, Andrei Butnaru, and Tommi Jauhiainen. 2019. A Report on the Third VarDial Evaluation Campaign. In *Proceedings of VarDial*.