

Domain Control for Neural Machine Translation

Catherine Kobus and Josep Crego and Jean Senellart

firstname.lastname@systrangroup.com

SYSTRAN International / 5 rue Feydeau, 75002 Paris, France

Abstract

Machine translation systems are very sensitive to the domains they were trained on. Several domain adaptation techniques have already been deeply studied. We propose a new technique for neural machine translation (NMT) that we call domain control which is performed at runtime using a unique neural network covering multiple domains. The presented approach shows quality improvements when compared to dedicated domains translating on any of the covered domains and even on out-of-domain data. In addition, model parameters do not need to be re-estimated for each domain, making this effective to real use cases. Evaluation is carried out on English-to-French translation for two different testing scenarios. We first consider the case where an end-user performs translations on a known domain. Secondly, we consider the scenario where the domain is not known and predicted at the sentence level before translating. Results show consistent accuracy improvements for both conditions.

1 Introduction

Machine translation systems are very sensitive to the domain(s) they were trained on because each domain has its own style, sentence structure and terminology. There is often a mismatch between the domain for which training data are available and the target domain of a machine translation system. If there is a strong deviation between training and testing data, translation quality will be dramatically deteriorated. Word ambiguities are often an issue for machine translation systems. For instance, the English word "*administer*" has to be

translated differently if it appears in medical or political contexts. Our work is motivated by the idea that neural models could benefit from having domain information to choose the most appropriate terminology and sentence structure while using the information from *all* the domains to improve the base translation quality. Recently, (Sennrich et al., 2016) report on the neural network ability to control politeness through side constraints. We extend this idea to domain control. Our goal is to allow a model built from a diverse set of training data to produce in-domain translations. This is, to extend the coverage of generic NMT models to specific domains, with their specialized terminology and style, without lowering translation quality on more generic data. We present two frameworks to feed domain meta-information on the NMT encoder side.

The paper is structured as follows: Section 2 overviews related work. Details of our neural MT engine are given in Section 3. Section 4 describes the proposed approach. Experiments and results are detailed in Section 5. Finally, conclusions and further work are drawn in Section 6.

2 Related Work

A lot of work has already been done for domain adaptation in Statistical Machine Translation. The approaches vary from in-domain data selection based methods (Hildebrand et al., 2005) (Moore and Lewis, 2010) (Sethy et al., 2006) to in-domain models mixture-based methods (Foster and Kuhn, 2007) (Koehn and Schroeder, 2007) (Schwenk and Koehn, 2008).

Recent works have especially dealt with domain adaptation for NMT by providing meta-information to the Neural Network. Our work is in line with this kind of approach. (Chen et al., 2016) feeds Neural Network with topic information on

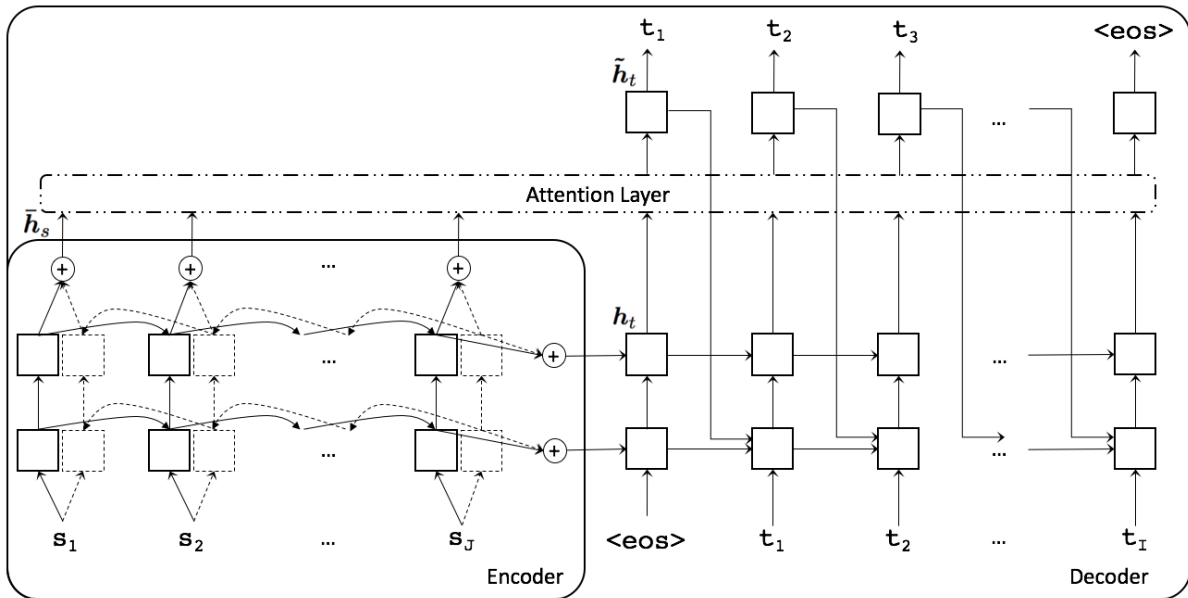


Figure 1: Schematic view of our NMT system.

the decoder side; topics are numerous and consist in human-labeled product categories. (Zhang et al., 2016) includes topic modeling on both encoder and decoder sides. A given number of topics are automatically inferred from the training data using Latent Dirichlet Allocation; each word in a sentence is assigned its own vector of topics. In our work, we also provide meta-information about domain to the network. However, we introduce domain information at the sentence level. (Luong and Manning, 2015) adapt a generic NMT network (trained on out-of-domain data) by running additional training iterations over an in-domain data set. The authors claim to obtain a domain adapted model in a very limited training time. However, it differs from our work since we aim at performing domain-adapted translations using a unique network that covers multiple domains.

3 Neural MT

Our NMT system follows the architecture presented in (Bahdanau et al., 2014). It is implemented as an encoder-decoder network with multiple layers of a RNN with Long Short-Term Memory hidden units (Zaremba et al., 2014). Figure 1 illustrates an schematic view of the MT network.

Source words are first mapped to word vectors and then fed into a bidirectional recurrent neural network (RNN) that reads an input sequence $s = (s_1, \dots, s_J)$. Upon seeing the $\langle \text{eos} \rangle$ symbol, the final time step initialises a target RNN.

The decoder is a RNN that predicts a target sequence $t = (t_1, \dots, t_I)$, being J and I respectively the source and target sentence lengths. Translation is finished when the decoder predicts the $\langle \text{eos} \rangle$ symbol.

The left-hand side of the figure illustrates the bidirectional encoder, which actually consists of two independent LSTM encoders: one encoding the normal sequence (solid lines) that calculates a forward sequence of hidden states $(\vec{h}_1, \dots, \vec{h}_J)$, the second encoder reads the input sequence in reversed order (dotted lines) and calculates the backward sequence $(\overleftarrow{h}_1, \dots, \overleftarrow{h}_J)$. The final encoder outputs $(\vec{h}_1, \dots, \vec{h}_J)$ consist of the sum of both encoders final outputs. The right-hand side of the figure illustrates the RNN decoder. Each word t_i is predicted based on a recurrent hidden state h_i and a context vector c_i that aims at capturing relevant source-side information.

Figure 2 illustrates the attention layer; it implements the "general" attentional architecture from (Luong et al., 2015). The idea of a global attentional model is to consider all the hidden states of the encoder when deriving the context vector c_t . Hence, global alignment weights a_t are derived by comparing the current target hidden state h_t with each source hidden state \vec{h}_s :

$$a_t(s) = \frac{\exp(\text{score}(h_t, \vec{h}_s))}{\sum_{s'} \exp(\text{score}(h_t, \vec{h}_{s'}))}$$

with the content-based score function:

$$\text{score}(h_t, \bar{h}_s) = h_t^T W_a \bar{h}_s$$

Given the alignment vector as weights, the context vector c_t is computed as the weighted average over all the source hidden states.




Figure 2: Attention layer of the NMT system.

The framework is available on the open-source project `seq2seq-attn`¹. More details about our system can be found in (Crego et al., 2016).

4 Domain control

Two different techniques are implemented to integrate domain control: additional token and domain feature.

4.1 Additional Token

The additional token method, inspired by the politeness control technique detailed in (Sennrich et al., 2016) consists in adding an artificial token at the end of each source sentence in order to let the network pay attention to the domain of each sentence pair. For instance, consider the next English-French translation:

Src: Headache may be experienced
Tgt: Des céphalées peuvent survenir

The network reads off the sentence pair with the appropriate *Medical* domain tag @MED@:

Src: Headache may be experienced @MED@
Tgt: Des céphalées peuvent survenir

¹<http://nlp.seas.harvard.edu>

Domain tags are appropriately selected in order to avoid overlaps with words present in the source language vocabulary. This method, though simple, has already proven to be effective to control the politeness level of a translation (Sennrich et al., 2016), or to support multi-lingual NMT models (Johnson et al., 2016).

4.2 Word Feature

We present a second technique to introduce domain control in our neural translation model. We use word-level features as described in (Crego et al., 2016). The first layer of the network is the word embedding layer. We adapt this layer to extend each word embedding with an arbitrary number of cells, designed to encode domain information. Notice that using additional features does not increase the vocabulary of source words; there are separate vocabularies for words and domain tags. Figure 3 illustrates a word embedding layer extended with domain information.




Figure 3: Word embedding layer for word w_j extended with domain label d , which constitutes a new input s_j for the encoder

Following with the example of Section 4.1, the sentence pair is given to the network with the appropriate *Medical* domain tag on each source word as follows:

Src: Headache may be experienced
MED MED MED MED
Tgt: Des céphalées peuvent survenir

Note that under this feature framework, the sentence-level domain information is added on a word-by-word basis to all the words in a sentence.

Domain	Lines	Src words	Tgt words	Lines	Src words	Tgt words
	Train			Test		
IT	399k	6.0M	7.3M	2k	36,8k	45,1k
Literature	35k	881k	943k	2k	50.1k	54.0k
Medical	923k	10.5M	12.3M	2k	35.6k	43.0k
News	194k	5.4M	6.7M	2k	53.5k	66.4k
Parliamentary	1.6M	37.6M	43.8M	2k	40.7k	49.4k
Tourism	1.1M	23.3M	27.5M	2k	39.1k	45.5k
Total	4,3M	83,7M	98.5M			

Table 1: Statistics for training and test sets of each domain corpus. Note that k stand for thousands and M for millions.

We reuse an existing framework that was originally implemented to include linguistic features at the word level (Crego et al., 2016).

5 Experiments

We evaluate the presented approach on English-to-French translation. Section 5.1 describes the data used for the experiments and details training configurations. Finally, Section 5.2 reports on translation accuracy results.

5.1 Training Details

We used training corpora covering six different domains: *IT*, *Literature*, *Medical*, *News*, *Parliamentary* and *Tourism*. *Medical*, *News*, *Parliamentary* data come from public corpora (respectively EMEA, News Commentary and Europarl), available from the OPUS repository (Tiedemann, 2012). *IT*, *Literature* and *Tourism* are proprietary data. Statistics of the corpora used are given in Table 1.

All experiments employ the NMT system detailed in Section 3 and are performed on NVidia GeForce GTX 1080. We use BPE² with a total of 32,000 source and target tokens as vocabulary, computed over the entire training corpora. Word embedding size is 500 cells. During training, we use stochastic gradient descent, a minibatch size of 64 with dropout probability set to 0.3 and bidirectional RNN. We train our models for 18 epochs. Learning rate is set to 1 and starts decaying after epoch 10 by 0.5. It takes about 10 days to train models on the complete training data set (4,3M sentence pairs).

Four different training configurations are considered. The first includes six in-domain NMT

models. Each model is trained using its corresponding domain data set (henceforth *Single* models). The *Join* network configuration is built using all the training data after concatenation. Note that this model does not include any information about domain. A *Token* network is also trained using all the available training data. It includes domain information through the additional token approach detailed in Section 4.1. Finally, *Feature* network is also trained on all available training data, it introduces domain information in the model by means of the feature framework detailed in Section 4.2.

5.2 Results

Table 2 shows translation accuracy results for the different training configurations. Accuracies are measured using BLEU³. As expected, the *Join* model outperforms all *Single* models on their corresponding test sets, showing that NMT engines benefit from additional training data. Differences in accuracy are lower for domains with a higher representation in the *Join* model, like *Parliamentary* and *Tourism*. No domain information is used on these first configurations (*none*).

Results for models incorporating domain information are detailed in columns *Token* and *Feature*. *Oracle* experiments indicate that the test set domains are known in advance, thus allowing to use the correct side-constraint. The additional token approach gives mixed results; it improves translation quality on some tasks and degrades on some others compared to the *Join* model. On the contrary, incorporating domain information through the *Feature* approach consistently improves translation quality on all the tasks. Adding domain information on all the source words seems to be a good technique to convey domain side-constraint

²<https://github.com/rsennrich/subword-nmt>

³multi-bleu.perl

Domain Constraint	Single None	Join	Token	Feature Oracle	Feature	
					RNN	Acc (%)
IT	52.73	53.81	53.76	54.56 (+0.75)	54.42	97.8
Literature	20.25	29.81	29.96	30.73 (+0.92)	30.71	93.1
Medical	33.97	41.83	42.02	42.51 (+0.68)	42.34	89.4
News	29.70	33.83	34.47	34.61 (+0.78)	34.49	88.3
Parliamentary	37.34	37.53	37.13	37.79 (+0.26)	37.77	82.7
Tourism	37.05	37.46	37.72	38.30 (+0.84)	38.01	90.6
Dialogs		19.25			19.55	

Table 2: BLEU scores for the different systems and RNN-based domain classifier accuracy.

Src:	Your doctor’s instructions should be carefully observed .
Ref:	Vous devrez respecter scrupuleusement les instructions de votre médecin .
Join:	Les instructions de votre médecin doivent être soigneusement surveillées .
Feature:	Les instructions de votre médecin doivent être suivies attentivement .
Src:	All injections of Macugen will be administered by your doctor.
Ref:	Toutes les injections de Macugen doivent être réalisées par votre médecin.
Join:	Toutes les injections de Macugen seront à l’ordre du jour de votre médecin.
Feature:	Toutes les injections de Macugen seront effectuées par votre médecin.

Table 3: Translation examples of in-domain medical sentences with and without domain feature

and to improve NMT target words choice consistency. Differences between the *Feature* and *Join* configurations are shown in parentheses. Note that an average improvement of 0.80 is observed on all test sets with the exception of Parliamentary translations, for which accuracy was only improved by 0.26. This can be explained by the fact that Parliamentary is the best represented domain in *Join* training set.

Translation examples are shown in Table 3 in a medical context. They show the impact on domain adaptation introduced by the *Feature* approach. The first example shows the preference of the *Feature* model for the French translation *suivies attentivement* of the English *carefully observed*. It seems more suitable than the hypothesis *soigneusement surveillées* output by the *Join*

model. A similar effect is shown on the second example where the French *effectuées* is clearly more adapted as translation of *administered* than *à l’ordre du jour*.

Finally, we also evaluate the ability of our presented approach (*Feature*) to face test sets for which the domain is not known in advance. Hence, before translation, the domain tag is automatically detected using an in-house domain classification module based on Recurrent Neural Networks (RNN) to disambiguate between the six different domains. The tool predicts the domain on a sentence-by-sentence basis, then translation is carried out using the predicted domain value in *Feature* model. Last column of Table 2 shows the accuracy of the domain classification tool for sentences on each of the predefined domains.

Test	Domain feature					
	IT	Literature	Medical	News	Parl.	Tourism
IT	54.56	-12.76	-10.25	-12.43	-13.83	-14.18
Literature	-5.96	30.73	-5.13	-2.89	-3.50	-3.03
Medical	-4.82	-6.23	42.51	-5.06	-5.39	-4.74
News	-3.36	-1.58	-3.04	34.61	-0.81	-2.48
Parliamentary	-4.14	-1.92	-3.09	-0.39	37.79	-3.01
Tourism	-6.72	-3.2	-4.16	-4.26	-4.35	38.30

Table 4: BLEU score decreases using different predefined domain tags

Results for this last condition are shown in column *RNN*. Even though domain is wrongly predicted in some cases, translation accuracy is still improved when compared to the *Join* model. Notice that domain classification at sentence level is a challenging task as short context is considered. We also confront our approach with a final test set from a brand new domain, *Dialogs*, that is not present in our training data. Sentences are selected from TED Talks corpora. The *RNN* toolkit is able to assign each test sentence to one of the source domains, leading to outperform the *Join* model.

In order to better understand the influence of the predicted domain, we conduct a final set of experiments. Using the *Feature* model, we run each test set using all domain values. Results are detailed in Table 4 showing that translation quality can significantly be degraded when translating sentences with the wrong domain tag. It is especially the case for *IT* domain, where translating with the wrong domain tag dramatically reduces accuracy. Results also reveal proximities between different domains like, for example, *News* and *Parliamentary*. Translating the *News* test set with the *Parliamentary* domain tag (and vice versa) does not seem to hurt translation quality compared to other domain tag mismatches.

6 Conclusions and Further Work

We have presented a method that incorporates domain information into a neural network. It allows to perform domain-adapted translations using a unique network that covers multiple domains. The presented method does not need to re-estimate model parameters when performing translations on any of the available domains.

We plan to further improve the feature technique detailed in this work. Rather than providing the network with a hard decision about domain, we want to introduce a vector of distance values of the given source sentence to each domain, thus allowing to smooth the proximity of each sentence to each domain.

Additionally, Table 4 shows indirectly that the neural network has learnt the ability to classify domains at the sentence level. We also plan to implement a joint approach for domain classification and translation, avoiding dependency with the RNN classifier.

Finally, since domain classification is a document level task, it would be interesting to extend

the current study to document level translation.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR* abs/1409.0473. Demoed at NIPS 2014: <http://lisa.iro.umontreal.ca/mt-demo/>.
- Wenhu Chen, Evgeny Matusov, Shahram Khadivi, and Jan-Thorsten Peter. 2016. Guided alignment training for topic-aware neural machine translation. *CoRR* abs/1607.01628v1.
- Josep Crego, Jungi Kim, Guillaume Klein, Anabel Rebollo, Kathy Yang, Jean Senellart, Egor Akhanov, Patrice Brunelle, Aurelien Coquard, Yongchao Deng, Satoshi Enoue, Chiyo Geiss, Joshua Johanson, Ardas Khalsa, Raoum Khiari, Byeongil Ko, Catherine Kobus, Jean Lorieux, Leidiana Martins, Dang-Chuan Nguyen, Alexandra Priori, Thomas Riccardi, Natalia Segal, Christophe Servan, Cyril Tiquet, Bo Wang, Jin Yang, Dakun Zhang, Jing Zhou, and Peter Zoldan. 2016. Systran’s pure neural machine translation systems. *CoRR* abs/1610.05540.
- George Foster and Roland Kuhn. 2007. Mixture-model adaptation for SMT. In *Proceedings of the Second Workshop on Statistical Machine Translation*. Association for Computational Linguistics, Prague, Czech Republic, pages 128–135.
- Almut Silja Hildebrand, Matthias Eck, Stephan Vogel, and Alex Waibel. 2005. Adaptation of the translation model for statistical machine translation based on information retrieval. In *Proceedings of the 10th Conference of the European Association for Machine Translation (EAMT)*. Budapest.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, and Niki Thorat. 2016. Google’s multilingual neural machine translation system: Enabling zero-shot translation arXiv:1611.04558v.
- Philipp Koehn and Josh Schroeder. 2007. Experiments in domain adaptation for statistical machine translation. In *Proceedings of the Second Workshop on Statistical Machine Translation*. Association for Computational Linguistics, Prague, Czech Republic, pages 224–227.
- Thang Luong and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *IWSLT2015*. Da Nang, Vietnam.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, pages 1412–1421.

- Robert C. Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 Conference Short Papers*. Association for Computational Linguistics, Uppsala, Sweden, pages 220–224.
- Holger Schwenk and Philipp Koehn. 2008. Large and diverse language models for statistical machine translation. In *Proceedings of the 3rd International Joint Conference on Natural Language Processing (IJCNLP)*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, USA, pages 35–40.
- Abhinav Sethy, Panayiotis Georgiou, and Shrikanth Narayanan. 2006. Selecting relevant text subsets from web-data for building topic specific language models. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*. Association for Computational Linguistics, New York City, USA, pages 145–148.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *LREC*. pages 2214–2218.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *CoRR* abs/1409.2329.
- Jian Zhang, Liangyou Li, Andy Way, and Qun Liu. 2016. Topic-informed neural machine translation. In *COLING*.