
Machine Translation in the Covid domain: an English-Irish case study for LoResMT 2021

Séamus Lankford seamus.lankford@adaptcentre.ie
ADAPT Centre, School of Computing, Dublin City University, Dublin, Ireland.

Haithem Affi haithem.affi@adaptcentre.ie
ADAPT Centre, Department of Computer Science, Munster Technological University, Ireland.

Andy Way andy.way@adaptcentre.ie
ADAPT Centre, School of Computing, Dublin City University, Dublin, Ireland.

Abstract

Translation models for the specific domain of translating Covid data from English to Irish were developed for the LoResMT 2021 shared task. Domain adaptation techniques, using a Covid-adapted generic 55k corpus from the Directorate General of Translation, were applied. Fine-tuning, mixed fine-tuning and combined dataset approaches were compared with models trained on an extended in-domain dataset. As part of this study, an English-Irish dataset of Covid related data, from the Health and Education domains, was developed. The highest-performing model used a Transformer architecture trained with an extended in-domain Covid dataset. In the context of this study, we have demonstrated that extending an 8k in-domain baseline dataset by just 5k lines improved the BLEU score by 27 points.

1 Introduction

Neural Machine Translation (NMT) has routinely outperformed Statistical Machine Translation (SMT) when large parallel datasets are available (Crego et al., 2016; Wu et al., 2016). Furthermore, Transformer based approaches have demonstrated impressive results in moderate low-resource scenarios (Lankford et al., 2021). NMT involving Transformer model development will improve the performance in specific domains of low-resource languages (Araabi and Monz, 2020). However, the benefits of NMT are less clear when using very low-resource Machine Translation (MT) on in-domain datasets of less than 10k lines.

The Irish language is a primary example of a low-resource language that will benefit from such research. This paper reports the results for the MT system developed for the English-Irish shared task at LoResMT 2021 (Ojha et al., 2021). Relevant work is presented in the background section followed by an overview of the proposed approach. The empirical findings are outlined in the results section. Finally, the key findings are presented and discussed.

2 Background

2.1 Transformer

A novel architecture called Transformer was introduced in the paper ‘Attention Is All You Need’ (Vaswani et al., 2017). Transformer is an architecture for transforming one sequence into another with the help of an Encoder and Decoder without relying on Recurrent Neural Networks.

Transformer models use attention to focus on previously generated tokens. This approach allows models to develop a long memory which is particularly useful in the domain of language translation.

2.2 Domain adaptation

Domain adaptation is a proven approach in addressing the paucity of data in low-resource settings. Fine-tuning an out-of-domain model by further training with in-domain data is effective in improving the performance of translation models (Freitag and Al-Onaizan, 2016; Sennrich et al., 2016). With this approach an NMT model is initially trained using a large out-of-domain corpus. Once fully converged, the out-of-domain model is further trained by fine-tuning its parameters with a low resource in-domain corpus.

A modification to this approach is known as mixed fine-tuning (Chu et al., 2017). With this technique, an NMT model is trained on out-of-domain data until fully converged. This serves as a base model which is further trained using the combined in-domain and out-of-domain datasets.

3 Proposed Approach

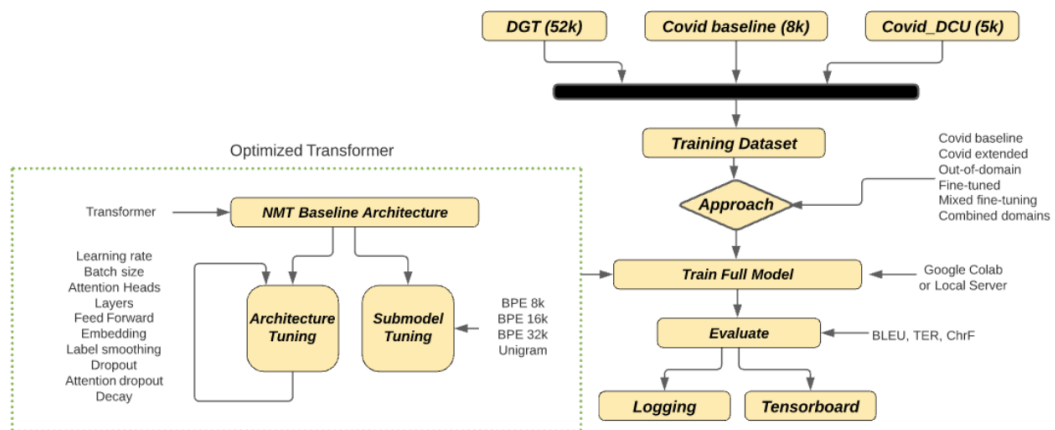


Figure 1: Proposed Approach. Optimal hyperparameters are applied to Transformer models which are trained using one of several possible approaches. The training dataset composition is determined by the chosen approach. Models are subsequently evaluated using a suite of metrics.

Hyperparameter optimization of Recurrent Neural Network (RNN) models in low-resource settings has previously demonstrated considerable performance improvements (Sennrich and Zhang, 2019). The extent to which such optimization techniques may be applied to Transformer models in similar low-resource scenarios was evaluated in a previous study (Lankford et al., 2021). Evaluations included modifying the number of attention heads, the number of layers and experimenting with regularization techniques such as dropout and label smoothing. Most importantly, the choice of subword model type and the vocabulary size were evaluated.

In order to test the effectiveness of our approach, models were trained using three English-Irish parallel datasets: a general corpus of 52k lines from the Directorate General for Translation (DGT) and two in-domain corpora of Covid data (8k and 5k lines). All experiments involved concatenating source and target corpora to create a shared vocabulary and a shared Sentence-Piece (Kudo and Richardson, 2018) subword model. The impact of using separate source and target subword models was not explored.

Approach	Source	Lines
Covid baseline	Baseline	8k
Covid extended	Baseline + Covid_DCU	13k
Out-of-domain	DGT	52k
Fine-tuned	Baseline + Covid_DCU + DGT	65k
Mixed fine-tuned	Baseline + Covid_DCU + DGT	65k
Combined domains	Baseline + Covid_DCU + DGT	65k

Table 1: Datasets used in proposed approach

Hyperparameter	Values
Learning rate	0.1, 0.01, 0.001, 2
Batch size	1024, 2048 , 4096, 8192
Attention heads	2 , 4, 8
Number of layers	5, 6
Feed-forward dimension	2048
Embedding dimension	128, 256 , 512
Label smoothing	0.1 , 0.3
Dropout	0.1, 0.3
Attention dropout	0.1
Average Decay	0, 0.0001

Table 2: Hyperparameter optimization for Transformer models. Optimal parameters are highlighted in bold (Lankford et al., 2021).

The approach adopted is illustrated in Figure 1 and the datasets used in evaluating this approach are outlined in Table 1. All models were developed using a Transformer architecture.

3.1 Architecture Tuning

Long training times associated with NMT make it costly to tune systems using conventional Grid Search approaches. A previous study identified the hyperparameters required for optimal performance (Lankford et al., 2021). Reducing the number of hidden layer neurons and increasing dropout led to significantly better performance. Furthermore, within the context of low-resource English to Irish translation, using a 16k BPE submodel resulted in the highest performing models. The Transformer hyperparameters, chosen in line with these findings, are outlined in Table 2.

4 Empirical Evaluation

4.1 Experimental Setup

4.1.1 Datasets

The performance of the Transformer approach is evaluated on English to Irish parallel datasets in the Covid domain. Three datasets were used in the evaluation of our models. These consisted of a baseline Covid dataset (8k) provided by MT Summit 2021 (Ojha et al., 2021), an in-domain Covid dataset (5k) developed at DCU and a publicly available out-of-domain dataset (52k) provided by DGT (Steinberger et al., 2013).




Figure 2: Translation performance of all approaches using Transformers with 2 heads

4.1.2 Infrastructure

Models were developed using a lab of machines each of which has an AMD Ryzen 7 2700X processor, 16 GB memory, a 256 SSD and an NVIDIA GeForce GTX 1080 Ti. Rapid prototype development was enabled through a Google Colab Pro subscription using NVIDIA Tesla P100 PCIe 16 GB graphic cards and up to 27GB of memory when available (Bisong, 2019).

Our MT models were trained using the Pytorch implementation of OpenNMT 2.0, an open-source toolkit for NMT (Klein et al., 2017).

4.1.3 Metrics

Automated metrics were used to determine the translation quality. All models were trained and evaluated using the BLEU (Papineni et al., 2002), TER (Snover et al., 2006) and ChrF (Popović, 2015) evaluation metrics. Case-insensitive BLEU scores, at the corpus level, are reported. Model training was stopped once an early stopping criteria of no improvement in validation accuracy for 4 consecutive iterations was recorded.

4.2 Results

Experimental results achieved using a Transformer architecture, with either 2 or 8 attention heads, are summarized in Table 3 and in Table 4. Clearly in the context of our low-resource experiments, it can be seen there is little performance difference using Transformer architectures with a differing number of attention heads. The largest difference occurs when using a fine-tuned approach (2.1 BLEU points). However the difference between a 2 head and an 8 head approach is less than 1 BLEU point for all other models. The highest performing approach uses the extended Covid dataset (13k) which is a combination of the MT summit Covid baseline and a custom DCU Covid dataset. This Transformer model, with 2 heads, performs well across all key translation metrics (BLEU: 36.0, TER: 0.63 and ChrF3: 0.32).

The worst performing model uses the Covid baseline which is not surprising given that only 8k lines are available. The performance of the higher resourced models (out-of-domain, fine-tuned, mixed fine-tuned and combined domains) all lag that of the Covid extended model. In particular, the out-of-domain model, using the DGT dataset, performs very poorly with a BLEU score of just 13.9 on a Transformer model with 2 heads.

The BLEU and TER scores for all approaches are illustrated in Figure 2a and Figure 2b. As expected, there is a high level of inverse correlation between BLEU and TER. Well-performing models, with high BLEU scores, also required little post editing effort as indicated by their lower TER scores.

System	Heads	Lines	Steps	BLEU \uparrow	TER \downarrow	ChrF3 \uparrow
Covid baseline	2	8k	35k	9.0	0.89	0.32
Covid extended	2	13k	35k	36.0	0.63	0.54
Out-of-domain	2	52k	200k	13.9	0.80	0.41
Fine-tuned	2	65k	35k	22.9	0.64	0.42
Mixed fine-tuned	2	65k	35k	18.2	0.71	0.42
Combined domains	2	65k	35k	32.2	0.59	0.55

Table 3: Comparison of optimized Transformer performance with 2 attention heads

System	Heads	Lines	Steps	BLEU \uparrow	TER \downarrow	ChrF3 \uparrow
Covid baseline	8	8k	35k	9.6	0.91	0.33
Covid extended	8	13k	35k	35.7	0.61	0.55
Out-of-domain	8	52k	200k	13.0	0.80	0.40
Fine-tuned	8	65k	35k	25.0	0.63	0.43
Mixed fine-tuned	8	65k	35k	18.0	0.71	0.42
Combined domains	8	65k	35k	32.8	0.59	0.57

Table 4: Comparison of optimized Transformer performance with 8 attention heads

5 Discussion

Standard Transformer parameters identified in a previous study were observed to perform well (Lankford et al., 2021). Reducing hidden neurons to 256 and increasing regularization dropout to 0.3 improved translation performance and these hyperparameters were chosen when building all Transformer models. Furthermore a batch size of 2048 and using 6 layers for the encoder / decoder were chosen throughout.

The results demonstrate that translation performance for specific domains is driven by the amount of data which is available for that specific domain. It is noteworthy that an in-domain dataset of 13k lines (Covid extended), trained for just 35k steps outperformed by 22.1 BLEU points the corresponding out-of-domain 52k dataset (DGT) which was trained for 200k steps.

6 Conclusion and Future Work

In the official evaluation for LoResMT 2021, our English–Irish system was ranked first according to the BLEU, TER and ChrF scores. We demonstrate that a high performing in-domain translation model can be built with a dataset of 13k lines. Developing a small in-domain dataset, of just 5k lines, improved the BLEU score by 27 points when models were trained with the combined Covid baseline and custom Covid dataset.

Following on from our previous work, careful selection of Transformer hyperparameters, and using a 16k BPE SentencePiece submodel, enabled rapid development of high performing translation models in a low-resource setting.

Within the context of our research in low-resource English to Irish translation, we have shown that augmenting in-domain data, by a small amount, performed better than approaches which incorporate fine-tuning, mixed fine-tuning or the combination of domains.

As part of our future work, we plan to develop English-Irish MT models trained on a dataset derived from the health domain. Domain adaptation, through fine-tuning such models with the Covid extended dataset may further improve Covid MT performance.

Acknowledgements

This work was supported by ADAPT, which is funded under the SFI Research Centres Programme (Grant 13/RC/2016) and is co-funded by the European Regional Development Fund. This research was also funded by the Munster Technological University.

References

- Araabi, A. and Monz, C. (2020). Optimizing transformer for low-resource neural machine translation. *arXiv preprint arXiv:2011.02266*.
- Bisong, E. (2019). Google colab. In *Building Machine Learning and Deep Learning Models on Google Cloud Platform*, pages 59–64. Springer.
- Chu, C., Dabre, R., and Kurohashi, S. (2017). An empirical comparison of domain adaptation methods for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 385–391.
- Crego, J., Kim, J., Klein, G., Rebollo, A., Yang, K., Senellart, J., Akhanov, E., Brunelle, P., Coquard, A., Deng, Y., et al. (2016). Systran’s pure neural machine translation systems. *arXiv preprint arXiv:1610.05540*.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., and Rush, A. M. (2017). Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*.
- Kudo, T. and Richardson, J. (2018). Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Lankford, S., Alfi, H., and Way, A. (2021). Transformers for low resource languages: Is feidir linn. In *Proceedings of the 18th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*.
- Ojha, A. K., Liu, C.-H., Kann, K., Ortega, J., Satam, S., and Fransen, T. (2021). Findings of the LoResMT 2021 Shared Task on COVID and Sign Language for Low-Resource Languages. In *Proceedings of the 4th Workshop on Technologies for MT of Low Resource Languages*.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (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.
- Popović, M. (2015). chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- Sennrich, R. and Zhang, B. (2019). Revisiting low-resource neural machine translation: A case study. *arXiv preprint arXiv:1905.11901*.
- Snover, M., Dorr, B., Schwartz, R., Micciulla, L., and Makhoul, J. (2006). A study of translation edit rate with targeted human annotation. In *Proceedings of association for machine translation in the Americas*, volume 200. Citeseer.
- Steinberger, R., Eisele, A., Kłoczek, S., Pilos, S., and Schlüter, P. (2013). Dgt-tm: A freely available translation memory in 22 languages. *arXiv preprint arXiv:1309.5226*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. *arXiv preprint arXiv:1706.03762*.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.