

Exploring Pair-Wise NMT for Indian Languages

by

Kartheek Akella, Sai Himallu, Sridhar Suresh Ragupathi, Aman Singhal, Zeeshan Khan, Vinay P
Namboodiri, C V Jawahar

in

Computation and Language

: 1

-7

Report No: IIIT/TR/2020/-1



Centre for Visual Information Technology
International Institute of Information Technology
Hyderabad - 500 032, INDIA
December 2020

Exploring Pair-Wise NMT for Indian Languages

Kartheek Akella*

CVIT, IIT-H

sukruthkartheek@gmail.com

Sai Himallu*

CVIT, IIT-H

saihimal.allu@gmail.com

Sridhar Suresh Ragupathi*

CVIT, IIT-H

srsridhar.98@gmail.com

Aman Singhal

CVIT, IIT-H

amansinghalml@gmail.com

Zeeshan Khan

CVIT, IIT-H

zeeshank606@gmail.com

Vinay P. Namboodiri

University of Bath

vpn22@bath.ac.uk

C V Jawahar

CVIT, IIT-H

jawahar@iiit.ac.in

Abstract

In this paper, we address the task of improving pair-wise machine translation for specific low resource Indian languages. Multilingual NMT models have demonstrated a reasonable amount of effectiveness on resource-poor languages. In this work, we show that the performance of these models can be significantly improved upon by using back-translation through a filtered back-translation process and subsequent fine-tuning on the limited pair-wise language corpora. The analysis in this paper suggests that this method can significantly improve a multilingual models' performance over its baseline, yielding state-of-the-art results for various Indian languages.

1 Introduction

Neural machine translation (NMT) algorithms, as is common for most deep learning techniques, work best with vast amounts of data. Various authors have argued that their performance would be limited for low resource languages Östling and Tiedemann (2017), Gu et al. (2018), Kim et al. (2020). One way to bridge this gap is through the use of multilingual NMT algorithms to bypass the data limitations of individual language pairs (Johnson et al. (2017), Aharoni et al. (2019), Vázquez et al. (2019)). The use of such a model has been demonstrated recently by Philip et al. (2021). In this paper, we investigate the problem of improving pair-wise NMT performance further over existing multilingual baselines. We specifically analyze the use of back-translation and fine-tuning to this effect. Our

*Equal Contribution. Sridhar worked on Tamil and Urdu, Himallu worked on Gujarati, Kartheek worked on Malayalam and Marathi, Aman worked on Hindi and Punjabi and Zeeshan worked on Odiya. Himallu was responsible for drafting this paper.

results suggest that it is possible to improve the performance of individual pairs of languages for various Indian languages. The performance of these language pairs is evaluated over standard datasets, and we observe consistent improvements over the two main baselines: an NMT system trained pair-wise from scratch using the corpora available for the pair of languages, a multilingual NMT model that uses many different languages.

2 Previous Work

The problem of multilingual NMT has attracted significant research attention in the recent past. (Dong et al., 2015) proposed the first multilingual model with a one-to-many mapping of languages, whereas (Ferreira et al., 2016) shared a single attention network for all language pairs. Recent works like (Conneau and Lample, 2019), (Conneau et al., 2020) which are an extension of (Liu et al., 2019) and (Devlin et al., 2019) have improved upon the initial formulation of the multilingual NMT problem. Works like (Currey et al., 2017), (Shah and Barber, 2018), (Li and Eisner, 2019) and (Hewitt and Liang, 2019) use monolingual data to supplement their parallel corpora to build an NMT system. On the flip side (Lample et al., 2018), (Wang et al., 2018), (Artetxe et al., 2018) study the unsupervised paradigm using only monolingual corpora.

Within the context of Indian languages, (Chandola and Mahalanobis, 1994) and (Dave et al., 2001) were one of the first works to explore a rule-based approach for translation from Hindi to English whereas (Patel et al., 2018), (Barman et al., 2014), (Saini and Sahula, 2018) and (Choudhary et al., 2018) have explored this problem through the prism of NMT. (Philip et al., 2019) and (Madaan and Sadat, 2020) extend the concept of multilingual

NMT to the setting of Indian languages. Due to the recent efforts undertaken by the authors of [Kakwani et al. \(2020\)](#), Indian languages are now better represented in terms of available monolingual corpora. These resources set up a fertile ground for exploitation by semi-supervised and unsupervised NMT approaches, which are consistent with the setting we study in this work.

3 Method

Consider a setting in which we have limited pair-wise corpora between a pair of languages, and we would like to obtain improved performance. We go about achieving this through the following procedure. First, we train a multilingual model on several languages. Next, we use existing monolingual corpora through back-translation (BT), and then we fine-tune the model using available pair-wise corpora. We show that this particular procedure indeed improves over the alternative approach of training a pair-wise NMT system using the available corpora. For the first step, we use the multilingual NMT model provided publicly by [\(Philip et al., 2021\)](#). We now provide details regarding the other two stages.

3.1 Back Translation

In NMT literature, BT is an effective approach that allows NMT to pivot from a fully supervised setting to a semi-supervised setting. When supplemented with other objectives (autoencoder denoising [\(Artetxe et al., 2018\)](#), cross-translation [\(García et al., 2020\)](#)), BT has demonstrated high efficacy in a fully unsupervised NMT setting as well. In the semi-supervised paradigm, which we study in this work, our object of interest is to generate a meaningful learning signal from monolingual resources of a particular language that allows a reasonable NMT model to exploit these resources to improve its performance on language pairs which include that specific language.

Filtering Mechanism

Using a reasonable NMT model, BT can leverage monolingual resources to generate notable amounts of low-quality synthetic parallel data. If low-quality parallel corpora can be filtered through some means such that erroneous translation pairs are eliminated, we can obtain a strong learning signal from such a filtered corpus. To design such a filtering mechanism, we draw inspiration from




Figure 1: SIM stands for a similarity heuristic. We use sentence wise BLEU scores in our work

generative modelling literature, precisely the idea of cyclic consistency [\(Zhu et al., 2017\)](#). Briefly, the idea of cyclical consistency within the context of computer vision relates to minimizing the discrepancy between an image from domain X and the image obtained after transforming it to a domain Y and then converting it back to the domain X. We adopt this approach to build our filtering mechanism, in which we first use our reasonable NMT model to generate intermediate English (EN) translations for the sentences in the monolingual corpus of some language (XX). As illustrated in fig 1, we then use these intermediate English translations to back-translate it into XX and then evaluate the sentence wise BLEU [\(Papineni et al., 2002\)](#) scores for each such translation as a measure of similarity. Only sentences that cross an empirically chosen threshold are retained to ensure that the generated translations are of good quality to obtain a reasonably high-quality synthetic parallel dataset. We refer to this filtering scheme as XX-EN-XX.

We initialise our NMT model using the weights from the multilingual NMT model provided by the authors of [\(Philip et al., 2021\)](#). The authors train a Transformer [\(Vaswani et al., 2017\)](#) model on 10 Indian languages namely (notation in brackets), Hindi (hi), Telugu (te), Tamil (ta), Malayalam (ml), Urdu (ur), Bangla (bn), Gujarati (gu), Marathi (mr) and Odia (od) in addition to English (en). They make two chief architectural decisions in this regard. One they develop a shared vocabulary over all languages of interest, giving equal representation to each language in the vocabulary (equal number of tokens from each language). The second, they share the encoder-decoder parameters of the Transformer model across all possible language pairs, a decision which encourages the model to learn a shared embedding space for all languages of interest. The

	Pre-filt #pairs	Post-filt #pairs
hi	4M	140K
pa	58K	7K
mr	178K	58K
gu	370K	39K
ta	88K	34K
ur	400K	105K
ml	178K	52K
od	221K	64K

Table 1: Monolingual corpora utilized

	hi	pa	gu	mr	ta	ml	ur	od
iitb	1.5M	-	-	-	-	-	-	-
cvit-pib	195K	27K	29K	81K	87K	32K	45K	-
ufal	-	-	-	-	167K	-	-	-
ilci	49K	49K	49K	-	49K	30K	49K	-
odcorp1.0	-	-	-	-	-	-	-	27K
odcorp2.0	-	-	-	-	-	-	-	97K
	1.75M	76K	79K	81K	303K	62K	94K	124K

Table 2: Parallel corpora utilized

low resource nature of these languages is primarily addressed through two techniques: namely, Transfer Learning and Backtranslation (Sennrich et al., 2016). This design choice allows us to use the same NMT model for the XX-EN and the EN-XX directions of the XX-EN-XX setting with a consistent amount of effectiveness. We reason that an EN-XX-EN filtering scheme would result in a compounding of errors problem due to the superior performances offered by the multilingual NMT model in the XX-EN direction in contrast to the EN-XX direction. In such a case, populating our filtered corpus would require selecting a lower value of the threshold, which would compromise the quality of translation pairs, thereby leading to a weak supervisory signal.

4 Experimental Setup

4.1 Training Details

We make use of the Transformer-Base, a part of Fairseq library (Ott et al., 2019), which is built with 6 encoder-decoder layers, each having 512 hidden units and a singular attention head as our NMT model. We initialise our model with the weights of the multilingual NMT model provided by the authors of (Philip et al., 2021). We also utilise the SentencePiece (Kudo and Richardson, 2018) models of (Philip et al., 2021) to build our vocabulary.

For all languages of interest, we carry out fil-

tering of the back-translated corpus by first evaluating the mean of sentence-wise BLEU scores for the cyclically generated translations and then selecting a value slightly higher than the mean as our threshold. Sentences that cross this threshold are then included along with their corresponding translations in our filtered corpus. We supplement the training of our NMT model on a filtered back-translated corpus with two rounds of finetuning on a relevant parallel corpus: a pre-training phase and a post-training phase before carrying out the final evaluation. We reason that a pre-training step enhances the possibility of generating a high-quality synthetic filtered corpus from the related monolingual corpora by providing a more robust prior NMT model for the BT routine. A post-training step ensures that the NMT model is subjected to a more reliable supervisory signal before the final evaluation is carried out. We train all our models using AdamW (Loshchilov and Hutter, 2019) optimizer until a local minimum is achieved.

4.2 Datasets

For Hi, Od, and Ta, we use the IIT-B corpus, ODIENCORP-v1.0, and PMIndia Corpus (Haddow and Kirefu, 2020) respectively. For the rest of the languages, we utilise the relevant monolingual corpora provided by the authors of (Kakwani et al., 2020). We use only a part of these monolingual corpora in our experiments, the statistics of which we present in Table 1. Our pre-training and post-training routines dictate the need for a parallel corpus for all our languages of interest. Due to the relative lack of availability of large high-quality parallel corpora for our language pairs of interest, we collate available resources to obtain a final parallel corpus, the details of which are presented in Table 2. In addition to CVIT-PIB (Siripragada et al., 2020) and ILCI (Jha, 2010) datasets, we also utilise IIT-B Hi-En corpus (Kunchukuttan et al., 2018) for Hi, UFAL ENTAMV2.0 (Ramasamy et al., 2012) for Ta and ODIENCORP 1.0 and 2.0 (Parida et al., 2020) for Od. For evaluation we use the CVIT-MKB (Siripragada et al., 2020) dataset for the languages in MKB. We evaluate on the ILCI dataset for Pa, OdiEncorp-v2.0 (Parida et al., 2020) for Od and UFAL EnTamv2.0 (Ramasamy et al., 2012) for Tamil.

PAIRS	Test-Set	State of the Art				NMT (Different Attempts)		
		Top-4 (Prev. Attempts)				Rand-init	M-NMT ¹	Filt-BT
En-Hi	MKB	15.65 ²	16.23 ²	21.05 ²	24.48²	13.28	16.93	16.67
En-Pa	ILCI				23.05 ¹	10.67	21.36	23.52
En-Mr	MKB	8.79 ²	8.84 ²	8.97 ²	9.65 ²	2.77	9.84	9.89
En-Gu	MKB	9.73 ²	10.13 ²	11.24 ²	11.70 ²	2.63	12.92	14.37
En-Ta	MKB	4.33 ²	4.43 ²	4.53 ²	4.94 ²	0.78	4.86	5.69
En-Ta	UFAL	11.73 ²	12.51 ²	12.74 ²	13.05 ²	0.78	7.80	19.07
En-MI	MKB	5.00 ²	5.17 ²	5.42 ²	6.32 ²	1.59	2.65	6.40
En-Ur	MKB				22.16 ¹	3.90	22.16	24.76
En-Od	ODIENCORPv2	7.93 ²	9.35 ²	9.85 ²	11.07²	5.29	0.96	10.84

Table 3: Comparison of our NMT results with others publicly available on WAT leader board ². For results that were not available on WAT leaderboard (Pa,Ur), we compare it with results from the paper (Philip et al., 2021). We find that initialisation using a multilingual model¹ is highly effective for NMT in contrast to initialising randomly and training only on the respective language

5 Results and Discussions

We report BLEU scores on all the test sets specified. We refer to our approach as Filt-BT in Table 3 and contrast our results with a randomly initialised model trained from scratch with the same conditions (Rand-Init), the multilingual model that we use as our prior NMT model (M-NMT) ¹ and the top 4 publicly available results on the WAT leaderboard ^{2 3}. Since Pa and Ur do not have an entry on the leaderboard, we instead make a comparison with the results reported in (Philip et al., 2021), which are the present SOTA results to the best of our knowledge.

The first comparison highlights the benefits of warm-starting our NMT model from a M-NMT model, whereas the second comparison helps us ascertain the efficacy of filtered BT and as such, we report consistent gains over both these baselines for all the language pairs. In all the language pairs barring Odia, we demonstrate the superior performances of a prior multilingual model in contrast to a specialized model trained from scratch, validating our claim that initialization using a multilingual model is highly effective for NMT in contrast to initializing randomly and training only on the respective language. Typically, we observe that using high threshold values for filtering leads to the filtered corpus getting biased by selecting comparatively shorter sentences. To maintain a healthy

mix of both types of sentences, we use a threshold value slightly higher than the mean of the sentence-wise BLEU scores which we find in our experiments empirically provides for a more balanced high quality (in terms of translation quality) corpus thereby guaranteeing a better supervisory signal.

For Ta and Ur, we notice a massive boosts in performance (11.88 and 10.49 BLEU points respectively) over our multilingual baseline, with significant gains, also being noticed for MI and Gu. The M-NMT model, which we use to initialize our NMT model, has been trained on OdiEnCorpV1.0, whereas our Rand-Init model has been trained using both versions of the dataset. We ascribe the former’s inferior performance compared to the latter on Odia to the domain mismatch between both these versions, something which the latter model does not have to face. We select a subset of the monolingual data to maintain consistency with our computing resources. For 200K sentences, we train on a 1080ti NVIDIA GPU and found that back translation took about 3 hours. We decided that it would be an adequate sample size to test the validity of our approach. Since only a subset of monolingual data provided by (Kakwani et al., 2020) is used, we fully expect these results to trend upwards if the entire corpora were to be utilised. We indicate the SOTA performance for each language in bold. As such, we achieve SOTA performances on Pa, Gu, MI, Mr, Ta and Ur.

6 Summary and Directions

Our explorations in the applicability of Neural Machine Translation for Indian languages lead to the

¹Philip et al. (2021)

²<http://lotus.kuee.kyoto-u.ac.jp/WAT/WAT2020/>

³We do not include results from model-ensembling approaches

following observations (i) Multilingual models are a promising direction to address data scarcity and the variability of resources across languages (ii) adapting a multilingual model for a specific pair can provide superior performances. We believe these solutions can further benefit from the availability of mono-lingual resources and noisy parallel corpora.

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised neural machine translation. In *Proceedings of the Sixth International Conference on Learning Representations*.
- Anup Barman, Jumi Sarmah, and Shikhar Sarma. 2014. [Assamese WordNet based quality enhancement of bilingual machine translation system](#). In *Proceedings of the Seventh Global Wordnet Conference*, pages 256–261, Tartu, Estonia. University of Tartu Press.
- Anoop Chandola and Abhijit Mahalanobis. 1994. Ordered rules for full sentence translation: A neural network realization and a case study for hindi and english. *Pattern Recognit.*, 27:515–521.
- Himanshu Choudhary, Aditya Kumar Pathak, Rajiv Ratan Saha, and Ponnurangam Kumaraguru. 2018. [Neural machine translation for English-Tamil](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 770–775, Belgium, Brussels. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 7059–7069. Curran Associates, Inc.
- Anna Currey, Antonio Valerio Miceli Barone, and Kenneth Heafield. 2017. [Copied monolingual data improves low-resource neural machine translation](#). In *Proceedings of the Second Conference on Machine Translation*, pages 148–156, Copenhagen, Denmark. Association for Computational Linguistics.
- Shachi Dave, Jignashu Parikh, and Pushpak Bhattacharyya. 2001. [Interlingua-based english-hindi machine translation and language divergence](#). *Machine Translation*, 16(4):251–304.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1723–1732, Beijing, China. Association for Computational Linguistics.
- Daniel C. Ferreira, André F. T. Martins, and Mariana S. C. Almeida. 2016. [Jointly learning to embed and predict with multiple languages](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2019–2028, Berlin, Germany. Association for Computational Linguistics.
- X. García, P. Forêt, Thibault Sellam, and Ankur P. Parikh. 2020. A multilingual view of unsupervised machine translation. *ArXiv*, abs/2002.02955.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor O.K. Li. 2018. [Universal neural machine translation for extremely low resource languages](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 344–354, New Orleans, Louisiana. Association for Computational Linguistics.
- Barry Haddow and Faheem Kirefu. 2020. [Pmindia – a collection of parallel corpora of languages of india](#).
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.

- Girish Nath Jha. 2010. [The TDIL program and the Indian language corpora initiative \(ILCI\)](#). In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Languages Resources Association (ELRA).
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhibo 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.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. IndicNLP Suite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages. In *Findings of EMNLP*.
- Yunsu Kim, Miguel Graça, and Hermann Ney. 2020. [When and why is unsupervised neural machine translation useless?](#) In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 35–44, Lisboa, Portugal. European Association for Machine Translation.
- Taku Kudo and J. Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *EMNLP*.
- Anoop Kunchukuttan, Pratik Mehta, and Pushpak Bhattacharyya. 2018. [The IIT Bombay English-Hindi parallel corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc' Aurelio Ranzato. 2018. Unsupervised machine translation using monolingual corpora only. In *International Conference on Learning Representations (ICLR)*.
- Xiang Lisa Li and Jason Eisner. 2019. [Specializing word embeddings \(for parsing\) by information bottleneck](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2744–2754, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#).
- Pulkit Madaan and Fatiha Sadat. 2020. [Multilingual neural machine translation involving Indian languages](#). In *Proceedings of the WILDRE5– 5th Workshop on Indian Language Data: Resources and Evaluation*, pages 29–32, Marseille, France. European Language Resources Association (ELRA).
- Robert Östling and Jörg Tiedemann. 2017. [Neural machine translation for low-resource languages](#).
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- 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.
- Shantipriya Parida, Satya Ranjan Dash, Ondřej Bojar, Petr Motlicek, Priyanka Pattnaik, and Debasish Kumar Mallick. 2020. [OdiEnCorp 2.0: Odia-English parallel corpus for machine translation](#). In *Proceedings of the WILDRE5– 5th Workshop on Indian Language Data: Resources and Evaluation*, pages 14–19, Marseille, France. European Language Resources Association (ELRA).
- Raj Patel, Prakash Pimpale, and Sasikumar Mukundan. 2018. [Machine translation in indian languages: Challenges and resolution](#). *Journal of Intelligent Systems*.
- Jerin Philip, Vinay P. Namboodiri, and C. V. Jawahar. 2019. [A baseline neural machine translation system for indian languages](#).
- Jerin Philip, Shashank Siripragada, Vinay P. Namboodiri, and C. V. Jawahar. 2021. Revisiting low resource status of indian languages in machine translation (pre-print available on arxiv). In *Proceedings of ACM India Joint International Conference on Data Science —& Management of Data*, Bangalore, India.
- Loganathan Ramasamy, Ondřej Bojar, and Zdeněk Žabokrtský. 2012. Morphological processing for english-tamil statistical machine translation. In *Proceedings of the Workshop on Machine Translation and Parsing in Indian Languages (MTPIL-2012)*, pages 113–122.
- S. Saini and V. Sahula. 2018. Neural machine translation for english to hindi. In *2018 Fourth International Conference on Information Retrieval and Knowledge Management (CAMP)*, pages 1–6.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages

86–96, Berlin, Germany. Association for Computational Linguistics.

Harshil Shah and David Barber. 2018. [Generative neural machine translation](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1346–1355. Curran Associates, Inc.

Shashank Siripragada, Jerin Philip, Vinay P. Namboodiri, and C V Jawahar. 2020. [A multilingual parallel corpora collection effort for Indian languages](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 3743–3751, Marseille, France. European Language Resources Association.

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*, pages 5998–6008.

Raúl Vázquez, Alessandro Raganato, Jörg Tiedemann, and Mathias Creutz. 2019. [Multilingual NMT with a language-independent attention bridge](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 33–39, Florence, Italy. Association for Computational Linguistics.

Wei Wang, Taro Watanabe, Macduff Hughes, Tetsuji Nakagawa, and Ciprian Chelba. 2018. [Denoising neural machine translation training with trusted data and online data selection](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 133–143, Belgium, Brussels. Association for Computational Linguistics.

Jun-Yan Zhu, T. Park, Phillip Isola, and A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251.