

# Neural Machine Translation for Low-Resourced Indian Languages

Himanshu Choudhary, Shivansh Rao, Rajesh Rohilla

Delhi Technological University (Formerly Delhi college of Engineering)  
himanshu.dce12@gmail.com, rao.shivansh570@gmail.com, rajesh@dce.ac.in

## Abstract

A large number of significant assets are available online in English, which is frequently translated into native languages to ease the information sharing among local people who are not much familiar with English. However, manual translation is a very tedious, costly, and time-taking process. To this end, machine translation is an effective approach to convert text to a different language without any human involvement. Neural machine translation (NMT) is one of the most proficient translation techniques amongst all existing machine translation systems. In this paper, we have applied NMT on two of the most morphological rich Indian languages, i.e. English-Tamil and English-Malayalam. We proposed a novel NMT model using Multihead self-attention along with pre-trained Byte-Pair-Encoded (BPE) and MultiBPE embeddings to develop an efficient translation system that overcomes the OOV (Out Of Vocabulary) problem for low resourced morphological rich Indian languages which do not have much translation available online. We also collected corpus from different sources, addressed the issues with these publicly available data and refined them for further uses. We used the BLEU score for evaluating our system performance. Experimental results and survey confirmed that our proposed translator (24.34 and 9.78 BLEU score) outperforms Google translator (9.40 and 5.94 BLEU score) respectively.

**Keywords:** Multihead self-attention, Byte-Pair-Encodding, MultiBPE, low-resourced, Morphology, Indian Languages

## 1. Introduction

Many populated countries such as India and China have several languages which change region by region. For example, India has 23 constitutionally recognized official languages (e.g., Hindi, Malayalam, Telugu, Tamil, and Punjabi) and numerous unofficial local languages. Not only big countries, even small countries also rich in language diversity. There are 851 languages spoken in Papua New Guinea, which is one of the smallest populated regions. In India, the population is about three billion, but only about 10% of them can speak English<sup>1</sup>. Some studies say that out of those 10% English speakers only 2% can talk, write, and examine English well, and rest 8% can merely recognize simple English and talk with a variety of accents. Thinking about a large number of valuable sources is available on the web in English and most people in India can not understand it well, it becomes important to translate such content into neighborhood languages to facilitate people. Sharing pieces of information between human beings is important not only for business purposes but also for sharing their emotions, reviews, and acts. For this, translation plays an essential role in minimizing the communication hole between different peoples. Considering the vast amount of text, it is not viable to translate them manually. Hence, it becomes crucial to translate text from one language (say, English) to other languages (say, Tamil, Malayalam) automatically. This technique is also referred to as *machine translation*.

English to Indian language translation poses the challenge of morphological and structural divergence. For instance, (i) the number of parallel corpora and (ii) differences between languages, mainly the morphological richness and variation in word order due to syntactical divergence. Indian languages (IL) suffers from both of these problems,

especially when they are being translated from English. Moreover, Indian languages such as Malayalam and Tamil differ not only in word order but are also more agglutinative as compared to English which is fusional. For instance, English has Subject-Verb-Object (SVO) whereas Tamil and Malayalam have Subject-Object-Verb (SOV). While syntactic differences contribute to difficulties of translation models, morphological differences contribute to data sparsity. We attempt to overcome both issues in this paper.

There are various papers on machine translation, but apart from foreign languages most of the works on Indian languages are limited to Hindi and on conventional machine translation techniques such as (Patel et al., 2018) and (Raju and Raju, 2016). Most of the previous work is focused on separating the words in suffix and prefix based on some rules and then applying translation techniques. We addressed this issue with BPE to make this whole process more efficient and reliable. Moreover, We observed that very less work is being done on low resourced Indian languages and techniques such as Byte-pair-encoding (BPEmb), MultiBPEmb, word-embedding, and self-attention are still unexplored which have shown a significant improvement in Natural Language Processing. Though unsupervised machine translation (Artetxe et al., 2017) is also in the focus of many researchers, still it is not as precise as supervised learning. We, also addressed that there is no trustworthy Public data available for the translation of such languages. Thus, in this paper, we have applied a neural machine translation technique with Multihead-self attention along with word embeddings and Pre-Trained Byte-Pair-Encoding. We worked on English-Tamil and English-Malayalam language pairs as it is one of the most difficult languages pair (ZdenekŽabokrtský, 2012) to translate due to morphological richness of Tamil and Malayalam language. A similar approach can be applied

<sup>1</sup><https://www.bbc.com/news/magazine-20500312>

to other languages as well. We obtained the data from En-Tamv2.0, Opus and UMC005, preprocessed them and evaluated our result using the evaluation metric BLEU. We used OpenNMT-py for the implementation of our models<sup>2</sup>. Experimental results, as well as the survey by native peoples, confirms that our result is far better than conventional translation techniques on Indian languages.

The Main contributions of our work are as follows:

- This is the first work to apply pre-trained BPE and MultiBPE embeddings on Indian language pairs (English-Tamil, English-Malayalam) along with Multihead self-attention technique.
- We achieved good accuracy with a relatively simpler model and in less training time rather than training on a complex neural network which requires much resources and time to train.
- We have addressed the issues with data preprocessing of Indian languages and shown why it is a crucial step in neural machine translation.
- We made our preprocessed data publicly available, which by our knowledge contains the largest number of a parallel corpus for the languages (English-Tamil, English-Malayalam, English-Telugu, English-Bengali, English-Urdu)
- Our model outperforms Google translator with a margin of 3.36 and an 18.07 BLEU score.

The paper is organized as follows. Sections Background and Approach describe the related work and the method that we used for our translator, respectively. Section experiments and Results show data preprocessing and results and analysis of our model. Finally, Section 5. concludes the paper and future work.

## 2. Background

A large amount of work has been reported on machine translation (MT) in the last few decades, the first one in the 1950s (Booth, 1955). Various approaches is used by researchers, such as rule-based (Ghosh et al., 2014), corpus-based (Wong et al., 2006), and hybrid-based approach (Salunkhe et al., 2016). Each approach has its own flaws and strength. Rule-based machine translation (RBMT) is MT systems based on the linguistic information about the source and target languages which is retrieved from ( multilingual, bilingual or monolingual) dictionaries and grammars covering the main syntactic, semantic and morphological regularities. It is further divided into transfer-based approach (TBA)(Shilon, 2011) and inter-lingual based approach (IBA). In the Corpus-based approach, we use a large-sized parallel corpus as raw data. This raw data contains ground truth translation for the desired languages. These corpora are used to train the model for translation. A corpus-based approach further classified in (i) statistical machine translation (SMT) (Patel et al., 2018) and (ii) example-based machine translation (EBMT) (Somers,

2003). SMT is the combination of decoding algorithms and basic statistical language models. EBMT, on the other hand, uses the translation examples and generates the new translation accordingly. It is done by finding the examples which are matching with the input. The alignment has to be performed after that to find out the parts of translation that can be reused. Hybrid-base machine translation combines any corpus-based approach and transfer approach in order to overcome their limitations. According to the recent research (Khan et al., 2017) the machine translation performance of Indian languages such as (*e.g.*, Hindi, Bengali, Tamil, Punjabi, Gujarati, and Urdu) is of an average of 10% accuracy. This demands the necessity of building better translation systems for Indian languages.

Unsupervised machine translation is further a new way of translation without using the parallel corpus, but the results are still not remarkable. On the other hand, NMT is an emerging technique and shown significant improvement in the translation results. In this paper (Hans and Milton, 2016) phrase-based hierarchical model is used and trained after morphological preprocessing. (Patel et al., 2017) trained their model after compound splitting and suffix separation. Many researchers also tried the same way and achieved a decent result on their respective datasets (Pathak and Pakray, ). We observed that morphological pre-processing, compound splitting and suffix or prefix separation can be overcome by using Byte-Pair-Encoding and produce similar or even better translation results without making the model complex.

## 3. Approach

In this paper, we present a neural machine translation technique using Multihead self-attention and word-embedding along with pre-trained Byte-Pair-Encoding (BPE) on our preprocessed dataset of Indian languages. We developed an efficient translation system, that overcomes the OOV (Out Of Vocabulary) and morphological analysis problem for Indian languages which do not have many translations available on the web. first, we provide an overview of NMT, Multi-head self-attention, word embedding, and Byte Pair Encoding. Next, we describe the framework of our translation model.

### 3.1. Neural Machine Translation Overview

Neural Machine translation is a powerful algorithm based on neural networks and uses the conditional probability of translated sentences to predict the target sentences of given source language (Revanuru et al., 2017a). When coupled with the power of attention mechanisms, this architecture can achieve impressive results with different variations. The following sub-sections provide an overview of basic sequence to sequence architecture, self-attention and other techniques that are used in our proposed translator.

#### 3.1.1. Sequence to sequence architecture

Sequence to sequence architecture is used for response generation whereas in Machine Translation systems it is used to find the relations between two language pairs. It consists of two important parts, an encoder, and a decoder. The encoder takes the input from the source language and the

<sup>2</sup><http://opennmt.net/OpenNMT-py/>




Figure 1: Seq2Seq architecture for English-Tamil

decoder leads to the output based on hidden layers and previously generated vectors. Let  $A$  be the source and  $B$  be a target sentence. The encoding part converts the source sentence  $a_1, a_2, a_3, \dots, a_n$  into the vector of fixed dimensions and the decoder part gives the word by word output using conditional probability. Here,  $A_1, A_2, \dots, A_M$  in the equation are the fixed size encoding vectors. Using chain rule, the Eq. 1 is transformed to the Eq. 2.

$$P(B/A) = P(B|A_1, A_2, A_3, \dots, A_M) \quad (1)$$

$$P(B|A) = P(b_i|b_0, b_1, b_2, \dots, b_{i-1}; a_1, a_2, a_3, \dots, a_m) \quad (2)$$

The decoder generates output using previously predicted word vectors and source sentence vectors in Eq. 1.

### 3.1.2. Attention Model

In a basic encoder-decoder architecture, encoder memorizes the whole sentence in terms of vector, and store it in the final activation layer, then the decoder uses that vector to generates the target sentence. This architecture works quite well for small sentences, but for larger sentences, maybe longer than 30 or 40 words, the performance degrades. To overcome this problem attention mechanisms play an important role. The basic idea behind this is that each time, when the model predicts an output word, it only uses the parts of input where the most relevant information is concentrated instead of the whole sentence. In other words, it only pays attention to some weighted words. Many types of attention mechanisms are used in order to improve the translation accuracy, but the multi-head self-attention overcomes most of the problems.

**Self-attention** In self-attention architecture (Vaswani et al., 2017) at every time step of an RNN, a weighted average of all the previous states will be used as an extra input to the function that computes the next state. With the self-attentive mechanism, the network can decide to attend to a state produced many time steps earlier. This means that the latest state does not need to store all the information. The mechanism also makes it easier for the gradient to flow more easily to all previous states, which can help against the vanishing gradient problem.

**Multi-Head Attention** When we have multiple queries  $q$ , we can combine them in a matrix  $Q$ . If we compute alignment using dot-product attention, the set of equations that are used to calculate context vectors can be reduced as shown in figure 3.  $Q, K$ , and  $V$  are mapped into lower-dimensional vector spaces using weight matrices and the results are used to compute attention (which we call a Head).




Figure 2: Attention model

In Multi-Head Attention we have  $h$  such sets of weight matrices which give us  $h$  Heads.

$$A(q, K, V) = \sum_i \frac{e^{q \cdot k_i}}{\sum_j e^{q \cdot k_j}} v_i$$

$$A(Q, K, V) = \text{softmax}(QK^T)V$$

Figure 3: Multi-Head Attention

### 3.1.3. Word Embedding

Word embedding is a unique way of representing the word in a vector space such that we can capture the semantic similarity of each word. Each word is represented in hundreds of dimensions. Generally, pre-trained embeddings are used trained on the larger data sets, and with the help of transfer learning, we convert the words from vocabulary to vector. (Cho et al., 2014).

### 3.1.4. Byte Pair Encoding

BPE (Gage, 1994) is a data compression technique that replaces the most frequent pair of bytes in a sequence. We use this algorithm for word segmentation, and by merging frequent pairs of characters or character sequences we can get the vocabulary of desired size (Sennrich et al., 2015). BPE helps in the suffix, prefix separation, and compound splitting which in our case used for creating new and complex words of Malayalam and Tamil language by interpreting them as sub-words units. We used BPE along with pre-trained fast-text word embeddings<sup>3</sup> (Heinzerling and Strube, 2018) for both the languages with the variation in the vocabulary size. In our model, we got the best results with vocabulary size 25000 and dimension 300.

**MultiBPEmb** MultiBPEmb is a collection of multiple languages subword segmentation models and pre-trained subword embeddings trained on Wikipedia data similar to monolingual BPE. On the contrary, instead of training one segmentation model for each language, here we train a single model and a single embedding for all the languages. We can also create a vocabulary of only two languages, source, and target. It deals with the mixed language sentences (Native language along with English) which are being popular nowadays on social media. Since our sentences were

<sup>3</sup><https://github.com/bheinzerling/bpemb>

ID	Language	Train	Test	Dev
1	Tamil	183451	2000	1000
2	Malayalam	548000	3660	3000
3	Telugu	75000	3897	3000
4	Bengali	658000	3255	3500
5	Urdu	36000	2454	2000

Table 1: Dataset for Indian Languages

clean it almost produced similar results, with variation in the BLEU score by 0.60 in Tamil and 1.15 in Malayalam.

## 4. Experimentation and Results

### 4.1. Evaluation Metric

BLEU score is a method to measure the difference between machine translation and human translation (Papineni et al., 2002). The approach works by matching n-grams in result translation to n-grams in the reference text, where unigram is a unique token, bigram is a word pair and so on. A perfect match results in a score of 1.0 or 100%.

### 4.2. Dataset

We obtained the data from different resources such as EnTamV2.0 (Ramasamy et al., 2012), Opus (Tiedemann, 2012) and UMC005 (Jawaid and Zeman, 2011). The sentences are of domain news, cinema, bible and movie subtitles. We combined and preprocessed the data of Tamil, Malayalam, Telugu, Bengali, and Urdu. After preprocessing (as described below) and cleaning, the dataset is split into train, test, and validation. Our final dataset is described in table 1. In our knowledge this is the largest, clean and preprocessed public dataset<sup>4</sup> available on the web for general purpose uses. As there is no publicly available dataset to compare various approaches on Indian languages, our datasets can be used to set baseline results to compare with.

### 4.3. Data Pre-processing

In the Research works (Hans and Milton, 2016) (Ramesh and Sankaranarayanan, 2018) EnTamV2.0 dataset is used. Also, the Opus dataset is a much widely used parallel corpus resource in various researcher’s works. However, we observed that in both of these well-known parallel resources there are many repeated sentences, which may results into the wrong results (can be higher or lower) after dividing into train, validation, and test sets, as many of the sentences, occur both in train and test sets. In most of the work, the focus relies on the models without interpreting the data which performs much better on our own test set rather than on general translated sentences. Thus, it is essential to analyses, correct and cleans the data before using it for the experiments. Researchers should also provide a detailed source of the corpus otherwise results can be misleading such as in paper (Revanuru et al., 2017b). We observed the following four important issues in the online available corpus.

- Sentence repetition with the same source and target.

- Different translations by the same source.
- Same translated sentences by different source sentences.
- Indian language tokenization.

To overcome the first issue, we took unique pairs from all the parallel sentences and removed the repeating ones. To tackle the second and third case we removed sentence pairs which were repeated more than twice and the difference between their length are in the window of 5 words. It is because for both of these cases we cannot identify that which source is correct for the same translation and which translated sentence is comes from the same source. We observed that there were some sentences, which were repeating even more than 20 times in the Opus dataset. This confuses the model to learn, identify and capture different features and overfits the model. Though data-augmentation (Fadaee et al., 2017) can improve the translation results, but in that case, the original data should be pre-processed, otherwise many augmented sentences may appear in both train and test data which leads to higher but wrong BLEU score as it will not work efficiently on new sentences.

For the tokenization of the English language, there are many libraries and frameworks such as (*e.g.*, Perl tokenizer) but these do not work well on the Indian languages, due to the difference between morphological symbols. The word-formation of Indian languages is quite different which we believed can only be handled by either special library for that particular language or by Byte-Pair-Encoding. In the case of BPE, we don’t need to tokenize the words which generally leads to better translation results.

After working on all these minor, but effective pre-processing we got our final dataset. While extracting the data from the web, we also removed sentences with a length greater than 50, known translated words in target sentences, noisy translations, and unwanted punctuations. For the reliability of data, we also took the help of native speakers of these languages.

### 4.4. Translator

We tried various new techniques as described above to get a better intuition of the effects on these two Indian language pairs. Our first model consists of 4 layer Bi-directional LSTM encoder and a decoder with 500 dimensions each along with a vocabulary size of 50,004 words for both source and target. First, we used Bahdanau’s attention and Adam optimizer with the dropout (regularization) of 0.3 and the learning rate 0.001. Here we used the 300 dimensional Pre-trained fast text<sup>5</sup> word embeddings for both the languages. Secondly, we used Pre-trained fast text Byte-Pair-Encoding<sup>6</sup> with the same attention. In the third model, we changed the attention to multi-head with 8 heads and 6 encoding and decoding layers. It shows an improvement of 1.2 and 6.18 BLEU scores for Tamil and Malayalam respectively. For the final model we used Multilingual fast text pre-trained Byte-pair-Encoddings<sup>7</sup> and got our final

<sup>5</sup><https://fasttext.cc/docs/en/crawl-vectors.html>

<sup>6</sup><https://github.com/bheinzerling/bpemb>

<sup>7</sup><https://nlp.h-its.org/bpemb/multi/>

<sup>4</sup><https://github.com/himanshudce/Indian-Language-Dataset>

best results of 9.67 and 25.36 respectively as shown in table 2 and table 3.

#### 4.5. Result

Our results is shown in table 2 and table 3. For Google translate we used Python API to translate the English sentences and compared the results with our various models. From the test results, It is observed that our model overcomes the OOV (Out of Vocabulary) problem in some cases, and is handy enough to be used in day to day life and official work.




Figure 4: English-Tamil model comparison with Google Translator Table2

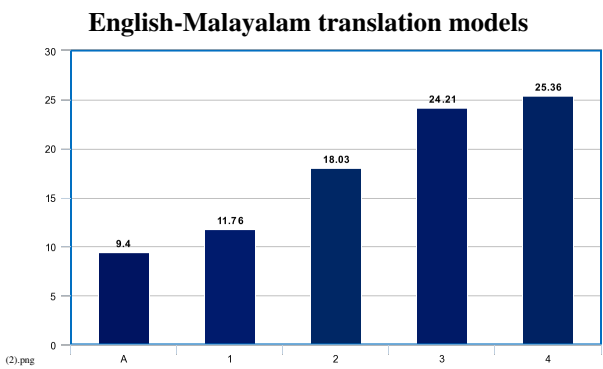


Figure 5: English-Malayalam model comparison with Google Translator Table 2

#### 4.6. Analysis

We conducted a survey with ten random sentences from our test data and accumulated the reviews of native Tamil speaking peoples. On comparing the reviews of Google translator and our translator, it was found, that our translation results were better in 60% cases than the Google translator. The visualization of an Attention can be seen in 6 of one of the sample sentences from our test data.

### 5. Conclusion

In this paper, we applied Neural Machine Translation (NMT) on two of the most difficult Indian language pairs (English-Tamil, English-Malayalam). We addressed the issues of data pre-processing and tokenization. To handle morphology and word complexities of Indian languages we

ID	Model	BLEU Score
A	Google Translator	5.71
1	Bi-LSTM(4-Layers)+ A+ Bahdanau Attention + WE	5.96
2	Bi-LSTM(4-Layers)+ A + Bahdanau Attention + Pre-BPE(25000)	7.87
3	Bi-LSTM(6-Layer)+ A + Multi-Head Attention+ Pre-BPE(25000)	9.07
4	Bi-LSTM(6-Layer)+ A + Multi-Head self Attention + Pre-MultiBPE(100000)	9.67

Table 2: English-Tamil model comparison with Google Translator ( A=Adam, WE=Word Embeddings)

ID	Model	BLEU Score
A	Google Translator	9.40
1	Bi-LSTM(4-Layers)+ A+ Bahdanau Attention + WE	11.76
2	Bi-LSTM(4-Layers)+ A + Bahdanau Attention + Pre-BPE(25000)	18.03
3	Bi-LSTM(6-Layer)+ A + Multi-Head Attention+ Pre-BPE(25000)	24.21
4	Bi-LSTM(6-Layer)+ A + Multi-Head self Attention + Pre-MultiBPE(100000)	25.36

Table 3: BLEU Score of English-Malayalam translated system. (A=Adam, B= Bahdanau, WE=Word Embedding)

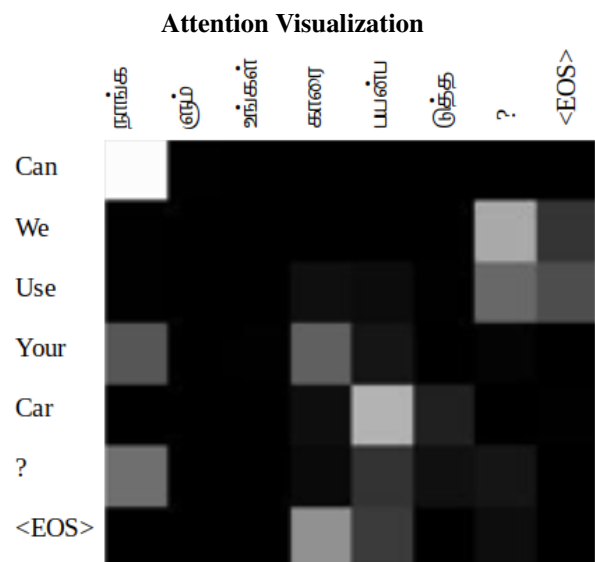


Figure 6: Attention visualization of English-Tamil sentence pair from our test data

applied pre-trained fast text BPEmb, MultiBPEmb embeddings along with multi-head self-attention which outperformed Google translator with a margin of 3.96 and 15.96 BLEU points respectively. The same approach can be ap-

plied to other Indian languages as well. Since the accuracy of our model was fairly good, it can be used for creating English-Malayalam and English-Tamil translation applications that will be very useful in domains like tourism, education and corporate. In the future, we can also explore the possibility to improve the translation results for code-switched languages using MultiBPE and other variations.

## 6. Bibliographical References

- Artetxe, M., Labaka, G., Agirre, E., and Cho, K. (2017). Unsupervised neural machine translation. *CoRR*, abs/1710.11041.
- Booth, A. D. (1955). Machine translation of languages, fourteen essays.
- Cho, K., Van Merriënboer, B., Bahdanau, D., and Bengio, Y. (2014). On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Fadaee, M., Bisazza, A., and Monz, C. (2017). Data augmentation for low-resource neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 567–573, Vancouver, Canada, July. Association for Computational Linguistics.
- Gage, P. (1994). A new algorithm for data compression. *The C Users Journal*, 12(2):23–38.
- Ghosh, S., Thamke, S., et al. (2014). Translation of telugu-marathi and vice-versa using rule based machine translation. *arXiv preprint arXiv:1406.3969*.
- Hans, K. and Milton, R. (2016). Improving the performance of neural machine translation involving morphologically rich languages. *arXiv preprint arXiv:1612.02482*.
- Heinzerling, B. and Strube, M. (2018). BPEmb: Tokenization-free Pre-trained Subword Embeddings in 275 Languages. In Nicoletta Calzolari (Conference chair), et al., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 7-12, 2018. European Language Resources Association (ELRA).
- Jawaid, B. and Zeman, D. (2011). Word-order issues in English-to-Urdu statistical machine translation.
- Khan, N. J., Anwar, W., and Durrani, N. (2017). Machine translation approaches and survey for indian languages. *arXiv preprint arXiv:1701.04290*.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Patel, R. N., Pimpale, P. B., et al. (2017). Mtil17: English to indian language statistical machine translation. *arXiv preprint arXiv:1708.07950*.
- Patel, R., Pimpale, P., and Mukundan, S. (2018). Machine translation in indian languages: Challenges and resolution. *Journal of Intelligent Systems*, 09.
- Pathak, A. and Pakray, P. ). Neural machine translation for indian languages. *Journal of Intelligent Systems*.
- Raju, B. N. V. N. and Raju, M. S. V. S. B. (2016). Statistical machine translation system for indian languages. In *2016 IEEE 6th International Conference on Advanced Computing (IACC)*, pages 174–177, Feb.
- Ramasamy, L., Bojar, O., and Žabokrtský, Z. (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.
- Ramesh, S. H. and Sankaranarayanan, K. P. (2018). Neural machine translation for low resource languages using bilingual lexicon induced from comparable corpora. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 112–119.
- Revanuru, K., Turlapaty, K., and Rao, S. (2017a). Neural machine translation of indian languages. In *Proceedings of the 10th Annual ACM India Compute Conference on ZZZ*, pages 11–20. ACM.
- Revanuru, K., Turlapaty, K., and Rao, S. (2017b). Neural machine translation of indian languages. 11.
- Salunkhe, P., Kadam, A. D., Joshi, S., Patil, S., Thakore, D., and Jadhav, S. (2016). Hybrid machine translation for english to marathi: A research evaluation in machine translation:(hybrid translator). In *Electrical, Electronics, and Optimization Techniques (ICEEOT), International Conference on*, pages 924–931. IEEE.
- Sennrich, R., Haddow, B., and Birch, A. (2015). Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Shilon, R. (2011). *Transfer-based Machine Translation between morphologically-rich and resource-poor languages: The case of Hebrew and Arabic*. Ph.D. thesis, Citeseer.
- Somers, H. (2003). An overview of ebmt. In *Recent advances in example-based machine translation*, pages 3–57. Springer.
- Tiedemann, J. (2012). Parallel data, tools and interfaces in opus. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, may. European Language Resources Association (ELRA).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Wong, F., Dong, M., and Hu, D. (2006). Machine translation using constraint-based synchronous grammar. *Tsinghua Science and Technology*, 11(3):295–306.
- ZdenekŽabokrtský, L. O. (2012). Morphological processing for english-tamil statistical machine translation. In *24th International Conference on Computational Linguistics*, page 113.