

Summarizing Indian Languages using Multilingual Transformers based Models

Dhaval Taunk¹, Vasudeva Varma¹

¹International Institute of Information Technology, Hyderabad, Telangana, India

Abstract

With the advent of multilingual models like mBART, mT5, IndicBART etc., summarization in low resource Indian languages is getting a lot of attention now a days. But still the number of datasets is low in number. In this work, we (Team HakunaMatata) study how these multilingual models perform on the datasets which have Indian languages as source and target text while performing summarization. We experimented with IndicBART and mT5 models to perform the experiments and report the ROUGE-1, ROUGE-2, ROUGE-3 and ROUGE-4 scores as a performance metric.

Keywords

Abstractive Summarization, mBART, mT5, IndicBART, ROUGE

1. Introduction

Automatic text summarization has a lot of potential applications in the current technological era like summarizing news articles, research articles etc. A lot of work has already been done in summarizing English languages text. But very little work is being done in summarizing Indian Languages. Therefore, summarizing text in these languages apart from English has become an essential task. India has approximately 350 million and 50 million Hindi and Gujarati speakers respectively. So building a summarization model in these languages will play a crucial role for this task. Recently, transformers based models like mBart[1], mT5[2] and IndicBart[3] have gained a lot of attention because of their multilingual capabilities including various Indic Languages.

Summarization can be performed in 2 ways: extractive summarization and abstractive summarization. In extractive summarization, a subset of sentences from the input text is taken as output summary. While in abstractive summarization, the entire summary is generated from scratch with the source text as input. Since text in abstractive summarization, summary is generated from scratch, this makes it more human like generated text. But at the same time, it becomes more difficult to perform abstractive summarization as compared to extractive summarization.

In this work, we aim to perform abstractive summarization on these languages as a part of the FIRE shared task 2022 - ILSUM [4][5] using the dataset provided by the organizers. We used

Forum for Information Retrieval Evaluation, December 9-13, 2022, India

✉ dhaval.taunk@research.iiit.ac.in (D. Taunk); vv@iiit.ac.in (V. Varma)

🌐 <https://dhavaltaunk08.github.io/> (D. Taunk); <https://www.iiit.ac.in/~vv> (V. Varma)

🆔 0000-0001-7144-4520 (D. Taunk)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

IndicBART and mT5 models for our experiments. We also performed data augmentation and tested the performance of the models. In the last, we report the ROUGE-1, ROUGE-2, ROUGE-3 and ROUGE-4 scores as mentioned by the shared task organizers.

2. Related Work

Both the extractive and abstractive summarization are well explored problem in English language context. A lot of datasets are available in English. Pubmed[6], arXiv[7], CNN/Daily Mail[8] are to name a few.

Guo et. al.,[9] extended T5[10] model to take long text as input and performed summarization over PubMed dataset. PRIMERA[11] is also another model which uses Longformer[12] model and achieved state of the art results on datasets like arXiv summarization data[7], Multi-News[13] and WCEP[14] datasets. Hasan et. al.,[15] introduced a multilingual dataset named XL-Sum comprising of 44 languages. They experimented with mT5[2] model to perform abstractive summarization and report results based upon that.

Aries et. al.,[16] performed multilingual and multi document summarization by clustering sentences into topics using a fuzzy clustering algorithm. They score each sentence based upon the topic coverage and then they create summary using the highest scoring sentences. For cross lingual abstractive summarization, Ladhak et. al.,[17] proposed WikiLingua, a article-summary pairs multilingual dataset available in 18 different languages. They fine-tuned mBART[1] in their experiments.

3. Methodology

The main aim of the task is to generate summary for the articles and headline pairs 3 languages viz. English, Hindi, Gujarati. Although news articles and headlines have been used in a number of earlier efforts in other languages, the current dataset presents a special problem of code- and script-mixing. Even though the item is written in an Indian language, phrases from English are frequently used in news stories. We perform experiments using IndicBart and mT5 models after performing some data analysis. We also found data augmentation to be a useful approach in getting better results.

3.1. Data Description

The dataset provided by the organizers is divided into 3 languages with 3 splits (train, val and test) present for all these 3 languages. Table 1 shows the article count for all 3 splits across all the languages.

For training phase, we were provided with the id of the article, link to the article, heading, summary and article text. While for the testing phase, we we given id of the article and the corresponding article text.

Since no references summaries was provided for the validation set. Therefore, while performing our own experiment, we took a small subset from the train set as our in-house validation set and then performed experiments over these sets. While we performed experiments on our

in-house validation set, during the validation phase, we evaluated our model on the official validation set. After that, since there was limitation of 3 submission per languages, we chose our top 3 performing experiments of each language as our final submission for test phase.

Table 1

Number of instances per languages per split

| Language | Train | Validation | Test |
|----------|-------|------------|------|
| English | 12565 | 898 | 4487 |
| Hindi | 7957 | 569 | 2842 |
| Gujarati | 8457 | 605 | 3020 |

4. Experiments

This section explains the steps that we undergo to perform the experiments and also the different experiments that we performed on the dataset.

4.1. Models used

For our experiments, we fine-tuned models viz. IndicBART and mT5-small the details of which is given below:

1. **IndicBART**: The eleven Indic languages and English are the main focus of the multilingual, sequence-to-sequence pre-trained model known as IndicBART. The authors tested IndicBART on two NLG tasks: extreme summarization and neural machine translation (NMT) and demonstrated that despite being substantially smaller, models IndicBART is competitive with huge pre-trained models like mBART50.
2. **mT5**: A new Common Crawl-based dataset with 101 languages was used to pre-train the multilingual T5 model (mT5 model). For mT5, the model design and training process closely resemble those of T5.

Both the models follows 12 layer (6 layer encoder + 6 layer decoder) architecture.

4.2. Data Augmentation

Apart from fine-tuning the model on actual test set, we also performed data augmentation and found a significant improvement over the results. For data augmentation, we performed 2 experiments. One by augmenting the 3X data to the actual dataset. Another by appending 5X data to the actual dataset. We found out that the performance of the model increased with increase in data augmentation amount.

4.3. Training Configuration

We used HuggingFace API and PyTorch to fine-tune the models. We used a learning rate of $2e-5$. Maximum input and output sequence length as 1024 and 100 respectively. And fine-tuned for 5,7, 10 epochs for different experiments.

5. Results

This section gives a detailed overview of results containing all the experiments performed on the validation set. While table 2, table 3 and table 4 gives results of various experiments on validation dataset for English, Hindi and Gujarati respectively. Along with that table 5, table 6 and table 7 shows the final 3 test set results per language.

5.1. Experiment Name

This subsection defines the experiment name with their details which are mentioned in the below mentioned tables:

5.1.1. English Experiments

1. **da_en_mt5**: mT5-small was finetuned in this approach along with data augmentation to 3 times of the actual english data.
2. **da_en_ibart**: IndicBART was finetuned in this approach along with data augmentation to 3 times of the actual english data.
3. **da5_en_ibart**: IndicBART was finetuned in this approach along with data augmentation to 5 times of the actual english data.
4. **en_ibart**: IndicBART was finetuned in this approach on the actual english dataset.
5. **en_mt5**: mt5-small was finetuned in this approach on the actual english dataset.

5.1.2. Hindi Experiments

1. **da5_hi_ibart**: IndicBART was finetuned in this approach along with data augmentation to 5 times of the actual hindi data.
2. **da_hi_ibart**: IndicBART was finetuned in this approach along with data augmentation to 3 times of the actual hindi data.
3. **da_hi_mt5**: mT5-small was finetuned in this approach along with data augmentation to 3 times of the actual hindi data.
4. **hi_ibart**: IndicBART was finetuned in this approach on the actual hindi dataset.
5. **hi_mt5**: mT5-small was finetuned in this approach on the actual hindi dataset.

5.1.3. Gujarati Experiments

1. **gu_ibart**: IndicBART was finetuned in this approach on the actual gujarati dataset.
2. **da_gu_ibart**: IndicBART was finetuned in this approach along with data augmentation to 3 times of the actual gujarati data.

3. **da5_gu_ibart**: IndicBART was finetuned in this approach along with data augmentation to 5 times of the actual gujarati data.
4. **gu_mt5**: mT5-small was finetuned in this approach on the actual gujarati dataset.

5.2. Validation set results

Below 3 tables shows results of our experiments on the validation set.

Table 2

ROUGE F1 scores on English Validation set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| da_en_mt5 | 0.54 | 0.43 | 0.41 | 0.40 |
| da_en_ibart | 0.51 | 0.38 | 0.36 | 0.35 |
| da5_en_ibart | 0.51 | 0.38 | 0.36 | 0.35 |
| en_ibart | 0.49 | 0.36 | 0.33 | 0.32 |
| en_mt5 | 0.47 | 0.34 | 0.32 | 0.31 |

Table 3

ROUGE F1 scores on Hindi Validation set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| da5_hi_ibart | 0.6104 | 0.515 | 0.488 | 0.475 |
| da_hi_ibart | 0.604 | 0.508 | 0.482 | 0.470 |
| da_hi_mt5 | 0.595 | 0.49 | 0.473 | 0.46 |
| hi_ibart | 0.594 | 0.497 | 0.471 | 0.458 |
| hi_mt5 | 0.54 | 0.438 | 0.412 | 0.398 |

Table 4

ROUGE F1 scores on Gujarati Validation set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| gu_ibart | 0.246 | 0.146 | 0.118 | 0.105 |
| da_gu_ibart | 0.239 | 0.144 | 0.118 | 0.105 |
| da5_gu_ibart | 0.235 | 0.137 | 0.11 | 0.096 |
| gu_mt5 | 0.206 | 0.114 | 0.09 | 0.079 |

5.3. Test set results

The below 3 tables shows the results of top 3 experiments per language on official test set.

Table 5

ROUGE F1 scores on English Test set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| da5_en_ibart | 0.521 | 0.401 | 0.378 | 0.369 |
| da_en_ibart | 0.512 | 0.389 | 0.366 | 0.358 |
| en_ibart | 0.493 | 0.367 | 0.344 | 0.336 |

Table 6

ROUGE F1 scores on Hindi Test set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| da5_hi_ibart | 0.592 | 0.491 | 0.464 | 0.451 |
| da_hi_ibart | 0.586 | 0.485 | 0.458 | 0.445 |
| hi_mt5 | 0.544 | 0.438 | 0.41 | 0.397 |

Table 7

ROUGE F1 scores on Gujarati Test set

| Experiment | ROUGE-1 | ROUGE-2 | ROUGE-3 | ROUGE-4 |
|--------------|---------|---------|---------|---------|
| da5_gu_ibart | 0.242 | 0.146 | 0.119 | 0.106 |
| da_gu_ibart | 0.241 | 0.145 | 0.120 | 0.107 |
| gu_mt5 | 0.203 | 0.115 | 0.094 | 0.084 |

5.4. Analysis

From the above results, we can say that data augmentation is a useful step as it has shown significant improvement of results over other experiments. Also, on comparing IndicBART and mT5, we can say that IndicBART performed better in most of the cases than mT5 for the summarization task. Further improvement can be made by using larger models like mbart-large or mt5-base/mt5-large models.

6. Conclusion

In this work, we presented our work for performing summarizing indian languages as part of the Forum for Information Retrieval Evaluation, 2022 shared task. We perform various experiments with multilingual transformer based models like IndicBART and mT5-small and achieved significant results. For Hindi and Gujarati languages, we stood at 2nd place. While for English language, we stood at 4th place. Due to computational constraints, we were not able to use larger models like mbart-large and mt5-base which could have performed even better. We hope this work will help future research in this direction.

References

- [1] Y. Liu, J. Gu, N. Goyal, X. Li, S. Edunov, M. Ghazvininejad, M. Lewis, L. Zettlemoyer, Multilingual denoising pre-training for neural machine translation, Transactions of the Association for Computational Linguistics 8 (2020) 726–742. URL: <https://aclanthology.org/2020.tacl-1.47>. doi:10.1162/tacl_a_00343.
- [2] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, C. Raffel, mT5: A massively multilingual pre-trained text-to-text transformer, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online, 2021, pp. 483–498. URL: <https://aclanthology.org/2021.naacl-main.41>. doi:10.18653/v1/2021.naacl-main.41.
- [3] R. Dabre, H. Shrotriya, A. Kunchukuttan, R. Puduppully, M. Khapra, P. Kumar, IndicBART: A pre-trained model for indic natural language generation, in: Findings of the Association for Computational Linguistics: ACL 2022, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 1849–1863. URL: <https://aclanthology.org/2022.findings-acl.145>. doi:10.18653/v1/2022.findings-acl.145.
- [4] S. Satapara, B. Modha, S. Modha, P. Mehta, Findings of the first shared task on indian language summarization (ilsum): Approaches, challenges and the path ahead, in: Working Notes of FIRE 2022 - Forum for Information Retrieval Evaluation, Kolkata, India, December 9-13, 2022, CEUR Workshop Proceedings, CEUR-WS.org, 2022.
- [5] S. Satapara, B. Modha, S. Modha, P. Mehta, Fire 2022 ilsum track: Indian language summarization, in: Proceedings of the 14th Forum for Information Retrieval Evaluation, ACM, 2022.
- [6] L. G. Galileo Mark Namata, Ben London, B. Huang, Query-driven active surveying for collective classification, in: International Workshop on Mining and Learning with Graphs, Edinburgh, Scotland, 2012.
- [7] C. B. Clement, M. Bierbaum, K. P. O’Keeffe, A. A. Alemi, On the use of arxiv as a dataset, 2019. URL: <https://arxiv.org/abs/1905.00075>. doi:10.48550/ARXIV.1905.00075.
- [8] A. See, P. J. Liu, C. D. Manning, Get to the point: Summarization with pointer-generator networks, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Vancouver,

- Canada, 2017, pp. 1073–1083. URL: <https://aclanthology.org/P17-1099>. doi:10.18653/v1/P17-1099.
- [9] M. Guo, J. Ainslie, D. Uthus, S. Ontanon, J. Ni, Y.-H. Sung, Y. Yang, LongT5: Efficient text-to-text transformer for long sequences, in: Findings of the Association for Computational Linguistics: NAACL 2022, Association for Computational Linguistics, Seattle, United States, 2022, pp. 724–736. URL: <https://aclanthology.org/2022.findings-naacl.55>. doi:10.18653/v1/2022.findings-naacl.55.
- [10] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *Journal of Machine Learning Research* 21 (2020) 1–67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [11] W. Xiao, I. Beltagy, G. Carenini, A. Cohan, PRIMERA: Pyramid-based masked sentence pre-training for multi-document summarization, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 5245–5263. URL: <https://aclanthology.org/2022.acl-long.360>. doi:10.18653/v1/2022.acl-long.360.
- [12] I. Beltagy, M. E. Peters, A. Cohan, Longformer: The long-document transformer, *arXiv:2004.05150* (2020).
- [13] A. Fabbri, I. Li, T. She, S. Li, D. Radev, Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 1074–1084. URL: <https://aclanthology.org/P19-1102>. doi:10.18653/v1/P19-1102.
- [14] D. Gholipour Ghalandari, C. Hokamp, N. T. Pham, J. Glover, G. Ifrim, A large-scale multi-document summarization dataset from the Wikipedia current events portal, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Online, 2020, pp. 1302–1308. URL: <https://aclanthology.org/2020.acl-main.120>. doi:10.18653/v1/2020.acl-main.120.
- [15] T. Hasan, A. Bhattacharjee, M. S. Islam, K. Mubashir, Y.-F. Li, Y.-B. Kang, M. S. Rahman, R. Shahriyar, XL-sum: Large-scale multilingual abstractive summarization for 44 languages, in: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, Association for Computational Linguistics, Online, 2021, pp. 4693–4703. URL: <https://aclanthology.org/2021.findings-acl.413>. doi:10.18653/v1/2021.findings-acl.413.
- [16] A. Aries, D. E. Zegour, K. W. Hidouci, AllSummarizer system at MultiLing 2015: Multilingual single and multi-document summarization, in: Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, Association for Computational Linguistics, Prague, Czech Republic, 2015, pp. 237–244. URL: <https://aclanthology.org/W15-4634>. doi:10.18653/v1/W15-4634.
- [17] F. Ladhak, E. Durmus, C. Cardie, K. McKeown, WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization, in: Findings of the Association for Computational Linguistics: EMNLP 2020, Association for Computational Linguistics, Online, 2020, pp. 4034–4048. URL: <https://aclanthology.org/2020.findings-emnlp.360>. doi:10.18653/v1/2020.findings-emnlp.360.