

LEVERAGING AUXILIARY DOMAIN PARALLEL DATA IN INTERMEDIATE TASK FINE-TUNING FOR LOW-RESOURCE TRANSLATION

Shravan Nayak^{‡,*}, Surangika Ranathunga[†], Sarubi Thillainathan[†], Rikki Hung[§], Anthony Rinaldi[§], Yining Wang[§], Jonah Mackey[§], Andrew Ho[§], En-Shiun Annie Lee[§]

[‡]Microsoft, [†]University of Moratuwa, [§]University of Toronto

ABSTRACT

NMT systems trained on Pre-trained Multilingual Sequence-Sequence (PMSS) models flounder when sufficient amounts of parallel data is not available for fine-tuning. This specifically holds for languages missing/under-represented in these models. The problem gets aggravated when the data comes from different domains. In this paper, we show that intermediate-task fine-tuning (ITFT) of PMSS models is extremely beneficial for domain-specific NMT, especially when target domain data is limited/unavailable and the considered languages are missing or under-represented in the PMSS model. We quantify the domain-specific results variations using a domain-divergence test, and show that ITFT can mitigate the impact of domain divergence to some extent.

1 INTRODUCTION

Pre-trained Multilingual Sequence-Sequence (PMSS) models such as mBART (Tang et al., 2021) and mT5 (Xue et al., 2021) have shown considerable promise over vanilla Transformer models for Neural Machine Translation (NMT). This promise persists to low-resource language translation (Thillainathan et al., 2021), which remains a challenge despite the recent advances in the field (Ranathunga et al., 2021).

In addition to the empirical analysis carried out during the introduction of these PMSS models (Tang et al., 2021), further empirical analysis for NMT was conducted by Wang et al. (2022); Liu et al. (2021a) and Lee et al. (2022). The latter two focused on low-resource language pairs and showed that results for languages unseen in the PMSS model are very low. Lee et al. (2022) attributed the results diversity to the amount of fine-tuning/pre-training data, noise in fine-tuning data, domain mismatch in training/testing data and language typology. Liu et al. (2021a) experimented with continuous pre-training (CPT) to include unseen languages into the model, but found that when the amount of parallel data used in the fine-tuning stage is very low, CPT is not effective, particularly for non-English-centric translations.

However, both Lee et al. (2022) and Liu et al. (2021a) considered only the case where the PMSS model is fine-tuned only once with a parallel dataset belonging to a particular domain. A look into the available corpora suggests that there are either noisy automatically created parallel corpora or manually curated small parallel corpora for hundreds of languages, across diverse domains (Tiedemann & Thottingal, 2020). Bapna et al. (2022) automatically mined bitext for over 1000 languages from the web. Artetxe et al. (2020) also point to several initiatives aimed at creating parallel resources at scale. This means, that for a given language pair, there can be several parallel datasets, belonging to different domains. In fact, Artetxe et al. (2020) argue that pure unsupervised NMT setup is not realistic given the availability of parallel data.

One way to utilize parallel data coming from auxiliary domains in building domain-specific NMT systems is Intermediate Task Fine-tuning (ITFT). In ITFT, the PMSS model is first fine-tuned with data from an auxiliary domain, and later fine-tuned on the target domain. The existing ITFT research for NMT considered only the case where the NMT model is tested only for a domain included in

*Correspondence to: Shravan Nayak <shravannayak.p@gmail.com>

fine-tuning (Verma et al., 2022; Adelani et al., 2022). Moreover, none of this research analysed the impact of domain divergence in ITFT. The other alternative to exploit auxiliary domain parallel corpora is to improve the pre-training stage of the PMSS model (Reid & Artetxe, 2021). However, they do not discuss the impact of the size and domain of the parallel data used during pre-training.

The objective of this research is to explore the benefit of ITFT for domain-specific low-resource language NMT (LRL-NMT). We experimented with several LRLs, where some are not in the selected PMSS model. We selected datasets from three domains and investigated the impact of ITFT on in-domain (train and test with the same domain data) and out-domain (train with a dataset of one domain and test with another) translation. The latter refers to zero-shot translation for the considered domain.

We show that ITFT is beneficial for domain-specific translation when the target domain data is unavailable/limited, specifically for languages missing/under-represented in the PMSS models. We also note a strong impact of divergence between the domains of final task and testing data. However, using a different domain for the ITFT dampens the impact of this domain divergence. Thus when multiple domain data is available for a language pair, consciously selecting data for ITFT is important. We also release a multi-way parallel bible dataset of 25k for the selected low-resource languages, which until now had less than 4k parallel sentences.

2 RELATED WORK

Before the PMSS era, researchers have experimented with Transfer Learning (TL) on vanilla Transformer (Vaswani et al., 2017) models and recurrent models. During TL, a low-resource translation task is trained on an NMT model, which has already been trained for a high-resource translation task (Lakew et al., 2018; Dabre et al., 2019; Maimaiti et al., 2020; Imankulova et al., 2019).

ITFT is based on the concept of TL, where a pre-trained language model is fine-tuned with an intermediate task, before the final task. ITFT has been extensively studied for encoder-based models such as BERT (Phang et al., 2018). Takeshita et al. (2022); Maurya et al. (2021), used ITFT on PMSS models for summarization and text generation, respectively. They investigated the use of auxiliary *tasks* in ITFT. As for NMT, Adelani et al. (2022) explored two variants: (1) fine-tune with data from an auxiliary domain and further fine-tune with the target domain and (2) fine-tune by combining data from different domains and further fine-tune with the target domain. However, they tested only for the in-domain setup. Verma et al. (2022) also experimented with data from one domain as the intermediate task, and with a multilingual intermediate task. In the latter, data comes from multiple languages but for the same domain. However, none of this research had a detailed look at the impact of domain difference of datasets.

3 EXPERIMENTAL SETTINGS

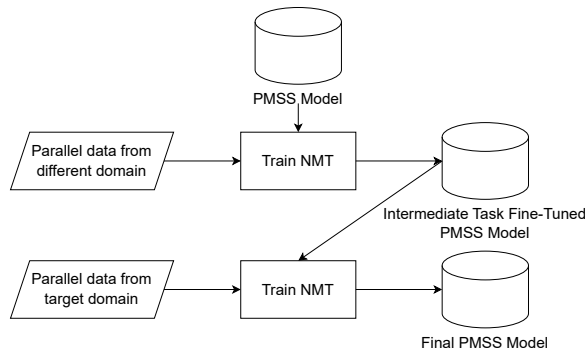


Figure 1: Overview of Multistage Fine-tuning.

ITFT: We fine-tune the PMSS model first with an out-domain parallel dataset (also referred to as *intermediate task*), followed by parallel data from another domain (also referred to as *final task*; see

Language	Family	Script	Joshi class	mBART Tokens (M)	mT5 Tokens (M)
Hindi (HI)	IA	Devanagari	4	1715	24000
Gujarati (GU)	IA	Gujarati	1	140	800
Kannada (KN)	Dr	Kannada	1	-	1100
Sinhala (SI)	IA	Sinhala	1	243	800
Tamil (TA)	Dr	Tamil	3	595	3400

Table 1: Languages (IA- Indo Aryan, Dr - Dravidian)

Dataset	Domain	Languages	Train Size	Test Size
FLORES-101	Open	HI, GU, KN, TA	test only	1k
FLORESv1	Open	SI	test only	1k
CCAligned (CC)	Open	ALL	100k	1k
Government (Gvt)	Administrative	SI, TA	50k	1k
PMIndia (PMI)	News	HI	50k	1k
		GU, KN	25k	1k
Web-scrap Bible	Religious	ALL	25k	1k

Table 2: Dataset statistics

Figure 1). Testing is carried out with data from the second domain (in-domain) or a different domain (out-domain).

Languages: We select four low-resource languages (Gujarati, Kannada, Sinhala and Tamil (Joshi et al., 2020)) along with English and Hindi. All, except English use non-Latin scripts (Pires et al., 2019). Table 1 reports further details.

Dataset: We use a mix of both open-domain and domain-specific corpora to train and test our models. The domain-specific corpora differ across the family of languages. Dataset summary details are given in Table 2. More details about the datasets and our new dataset can be found in Appendix A.1. Note that PMI and Government (Gvt) corpora are mutually exclusive for the datasets we considered¹. Therefore, when describing results (Section 4), we use PMI/Gvt to denote that we use one of these corpora for the considered experiment.

PMSS Models: Past work reported mixed results on the comparative performance of the two commonly used PMSS models, mBART and mT5 (Lee et al., 2022; Liu et al., 2021b). Thus we considered both these models for initial experiments. Comprehensive training details are given in Appendix A.2.

4 RESULTS AND DISCUSSION

We use SentencePiece BLEU (spBLEU) (Goyal et al., 2022) as the evaluation metric. Reasons for selecting this metric are discussed in Appendix A.3. For all our experiments, we obtain results for XX-EN, as well as EN-XX tasks. Observations discussed in this section hold for both translation directions, unless specifically mentioned. We carry out both *out-domain testing* as well as *in-domain testing*. Test set specifications are as indicated in Table 2.

4.1 BASELINE RESULTS

We fine-tune mBART and mT5 with each of the training sets, and evaluate with the test set. As per Figure 4 in Appendix B, mBART outperforms mT5 across domains and dataset sizes (except for Kannada, which is missing in mBART), for both in-domain and out-domain testing, and confirm Lee et al. (2022); Liu et al. (2021b)’s observations. Thus we selected mBART for further experiments.

Our mBART results in Table 5 in Appendix B echo Lee et al. (2022)’s observations: For the in-domain cases, NMT models built on mBART produce very low results for Kannada, when the par-

¹Although Tamil data is available in the PMI corpus we do not use this for our experiments.






Figure 2: In-domain testing with ITFT. Intermediate task - CC, final task - PMI and test - PMI. X-axis shows the final task size. Colours show the intermediate task size (size 0 = baseline). Train Size 1 corresponds to size of CC data set and similarly Train Size 2 for PMI/Gvt.



(a) Intermediate Task Fine-Tuning with Intermediate task - CC, final task PMI. The spBLEU scores corresponds to scores on FLORES test set.



(b) Intermediate Task Fine-Tuning with Intermediate task - CC, final task Bible. The spBLEU scores corresponds to scores on FLORES test set.

Figure 3: Out-domain testing with ITFT. intermediate task - CC, test- FLORES, (a) final task - PMI, (b) final task - Bible. X-axis shows the final task size. Colours show the intermediate task size (size 0 = baseline). Train Size 1 corresponds to size of the intermediate task data set and similarly Train Size 2 for final task.

allel data size is less than 10k. However, with 25k parallel sentences, even for Kannada, the model reports very strong results. This strong result confirms the data efficiency of the models trained on mBART. However, out-domain results are beyond usable levels for most of the cases. We analyse this observation with respect to domain divergence in Section 4.3.

4.2 EFFECTIVENESS OF ITFT

We vary the size and domain of the intermediate task, as well as the size of the final task.

In-domain translation: Here, the domain of the final task is similar to that of the testing data. We considered PMI/Gov't and Bible domains for the final task, since they have data for both training and testing. Figure 2 shows the result when the intermediate task is CC and final task is PMI. The complete result is reported in Figure 6 in Appendix B.

We see that ITFT with a different domain outperforms the baseline when the final task has less than 10k data (except for the 1k size of intermediate task for Ta and Si, where results are on par with the baseline). More training data for the intermediate task always helps in this setup. However, when final task dataset size increases, the impact of intermediate task saturates and finally converges to the baseline result. The language-wise observation is worth noting - the impact of ITFT is the highest for Kannada. The gains over the baseline are visible even when the final task dataset size reaches 50k. However, gain of ITFT is the least for Hindi, which is relatively highly represented in mBART.

Dataset	Gvt test	FLORES test	Bib test	PMI test
Gvt train	0.10	0.40	0.58	-
CC train	0.40	0.41	0.52	0.44
Bib train	0.56	0.47	0.10	0.53
PMI train	-	0.33	0.52	0.15

Table 3: JS divergence between train and test domains

Intermediate-final task size	Intermediate task	Final task
1k - 1k	0.07	0.16
25k - 25k	0.1	0.87
1k - 25k	0.49	0.95
25k - 1k	0.25	0.04

Table 4: Correlation between the out-domain result and the divergence between train/test domains

Out-domain translation: In Figure 3 we show the results for the cases where the intermediate task is CC, the final task is PMI/Gov’t and Bible, respectively, and the testing task is FLORES. Results for all the domain combinations are shown in Figure 6 in Appendix B.

Similar to the in-domain setup, we see that ITFT is very beneficial when the final task has less than 10k parallel data, and having more data for the intermediate task is always beneficial. However, when the intermediate task has only 1k parallel sentences, its effect is on par with that of the baseline for some cases (e.g. Kannada: intermediate - Bible, final - PMI). Different to in-domain translation, ITFT noticeably outperforms the baseline even for Hindi. However, we see that when Bible is used as the final task, the baseline shows very small gains as the final task size is increased. The gain slowly diminishes as the final task size is increased and finally tends to converge with the baseline result.

4.3 IMPACT OF DOMAIN DIVERGENCE

Lee et al. (2022) observed an impact of domain divergence in their results (i.e. our baseline), however did not provide any quantitative analysis. In contrast, we systematically quantify the impact of domain difference.

We measured the similarity between the two domains using the Jenson-Shannon (JS) divergence as defined in Plank & van Noord (2011), which is a modification of the Kullback-Leibler (KL) divergence (See Appendix A.4 for details of the calculation). Divergence between the training and test sets we used is given in Table 3.

To begin with, correlation between baseline results (Table 5) and the domain divergence values (Table 3) is $R^2 = 0.95$ - a very strong correlation. For example, when FLORES is the test set, fine-tuning with PMI/Gov’t gives much higher results than when using Bible, which is explained by the higher divergence between FLORES and Bible.

In order to quantify the interplay between size and domain of intermediate task fine-tuning, we calculated the correlation values as shown in Table 4. When the size of the intermediate task is 25k, the divergence between the domain of the final task and the domain of the test task shows a lesser correlation to that of the baseline. This means ITFT slightly dampens the effect of domain divergence of the final task, although the domain of the intermediate task itself does not have a noticeable impact on the result. When the intermediate task size is 1k, it makes no impact on the divergence between the final and test task domains. This could be the reason why the 1k ITFT result lies closer to the baseline result. We also note that when the final task size is 1k, the correlation is not well-defined, despite the size of intermediate task.

5 CONCLUSION

We analysed the possibility of using auxiliary parallel data in an ITFT setup to improve the PMSS model performance for low-resource domain-specific NMT. Our results show that ITFT is beneficial when limited data is available for a particular domain. This observation strongly holds for languages missing or under-represented in the PMSS model. We also show that the domain differences play a major role in out-domain translation, however ITFT can mitigate this impact to a certain extent. In future, we expect to combine parallel data from different domains in a multi-task setup to further mitigate the impact domain divergence.

REFERENCES

- David Ifeoluwa Adelani, Jesujoba Oluwadara Alabi, Angela Fan, Julia Kreutzer, Xiaoyu Shen, Machel Reid, Dana Ruitter, Dietrich Klakow, Peter Nabende, Ernie Chang, et al. A few thousand translations go a long way! leveraging pre-trained models for african news translation. *arXiv preprint arXiv:2205.02022*, 2022.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. A call for more rigor in unsupervised cross-lingual learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7375–7388, 2020.
- Ankur Bapna, Isaac Caswell, Julia Kreutzer, Orhan Firat, Daan van Esch, Aditya Siddhant, Mengmeng Niu, Pallavi Baljekar, Xavier Garcia, Wolfgang Macherey, et al. Building machine translation systems for the next thousand languages. *arXiv preprint arXiv:2205.03983*, 2022.
- Raj Dabre, Atsushi Fujita, and Chenhui Chu. Exploiting multilingualism through multistage fine-tuning for low-resource neural machine translation. 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)*, pp. 1410–1416, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1146. URL <https://aclanthology.org/D19-1146>.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. CCAIghned: A massive collection of cross-lingual web-document pairs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5960–5969, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.480. URL <https://aclanthology.org/2020.emnlp-main.480>.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. Beyond english-centric multilingual machine translation, 2020.
- Aloka Fernando, Surangika Ranathunga, and Gihan Dias. Data augmentation and terminology integration for domain-specific Sinhala-English-Tamil statistical machine translation. *arXiv preprint arXiv:2011.02821*, 2020. URL <https://arxiv.org/abs/2011.02821>.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzman, and Angela Fan. The FLORES-101 evaluation benchmark for low-resource and multilingual machine translation. *arXiv preprint arXiv:2106.03193*, 2021. URL <https://arxiv.org/abs/2106.03193>.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. The Flores-101 Evaluation Benchmark for Low-Resource and Multilingual Machine Translation. *Transactions of the Association for Computational Linguistics*, 10:522–538, 05 2022. ISSN 2307-387X. doi: 10.1162/tacl_a_00474. URL https://doi.org/10.1162/tacl_a_00474.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English. 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)*, pp. 6098–6111, Hong Kong, China, 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1632. URL <https://aclanthology.org/D19-1632>.
- Barry Haddow and Faheem Kirefu. PMIndia - A collection of parallel corpora of languages of India. *CoRR*, abs/2001.09907, 2020. URL <https://arxiv.org/abs/2001.09907>.
- Aizhan Imankulova, Raj Dabre, Atsushi Fujita, and Kenji Imamura. Exploiting out-of-domain parallel data through multilingual transfer learning for low-resource neural machine translation. In *Proceedings of Machine Translation Summit XVII: Research Track*, pp. 128–139, Dublin, Ireland, August 2019. European Association for Machine Translation. URL <https://aclanthology.org/W19-6613>.

- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6282–6293, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.560. URL <https://aclanthology.org/2020.acl-main.560>.
- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmungkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroko Orife, Kelechi Ogueji, Andre Niyongabo Rubungo, Toan Q. Nguyen, Mathias Müller, André Müller, Shamsuddeen Hassan Muhammad, Nanda Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, Mathias Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine Çabuk Ballı, Stella Biderman, Alessia Battisti, Ahmed Baruwā, Ankur Bapna, Pallavi Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. Quality at a Glance: An Audit of Web-Crawled Multilingual Datasets. *Transactions of the Association for Computational Linguistics*, 10:50–72, 01 2022. ISSN 2307-387X. doi: 10.1162/tacl.a.00447. URL <https://doi.org/10.1162/tacl.a.00447>.
- Surafel M Lakew, Aliia Erofeeva, Matteo Negri, Marcello Federico, and Marco Turchi. Transfer learning in multilingual neural machine translation with dynamic vocabulary. *arXiv preprint arXiv:1811.01137*, 2018.
- En-Shiun Lee, Sarubi Thillainathan, Shravan Nayak, Surangika Ranathunga, David Adelani, Ruisi Su, and Arya D McCarthy. Pre-trained multilingual sequence-to-sequence models: A hope for low-resource language translation? In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 58–67, 2022.
- Zihan Liu, Genta I Winata, Samuel Cahyawijaya, Andrea Madotto, Zhaojiang Lin, and Pascale Fung. On the importance of word order information in cross-lingual sequence labeling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13461–13469, May 2021a. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17588>.
- Zihan Liu, Genta Indra Winata, and Pascale Fung. Continual mixed-language pre-training for extremely low-resource neural machine translation. *arXiv preprint arXiv:2105.03953*, 2021b.
- Mieradilijiang Maimaiti, Yang Liu, Huanbo Luan, and Maosong Sun. Enriching the transfer learning with pre-trained lexicon embedding for low-resource neural machine translation. *Tsinghua Science and Technology*, pp. 1, 2020.
- Kaushal Kumar Maurya, Maunendra Sankar Desarkar, Yoshinobu Kano, and Kumari Deepshikha. Zmbart: An unsupervised cross-lingual transfer framework for language generation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 2804–2818, 2021.
- Arya D. McCarthy, Rachel Wicks, Dylan Lewis, Aaron Mueller, Winston Wu, Oliver Adams, Garrett Nicolai, Matt Post, and David Yarowsky. The Johns Hopkins University Bible corpus: 1600+ tongues for typological exploration. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 2884–2892, Marseille, France, 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.352>.
- Jason Phang, Thibault Févry, and Samuel R Bowman. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*, 2018.
- Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual BERT? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4996–5001, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1493. URL <https://aclanthology.org/P19-1493>.

- Barbara Plank and Gertjan van Noord. Effective measures of domain similarity for parsing. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 1566–1576, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://aclanthology.org/P11-1157>.
- Maja Popović. chrF++: words helping character n-grams. In *Proceedings of the Second Conference on Machine Translation*, pp. 612–618, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-4770. URL <https://aclanthology.org/W17-4770>.
- Matt Post. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 186–191, Brussels, Belgium, October 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-6319. URL <https://aclanthology.org/W18-6319>.
- Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. Neural machine translation for low-resource languages: A survey. *arXiv preprint arXiv:2106.15115*, 2021.
- Machel Reid and Mikel Artetxe. Paradise: Exploiting parallel data for multilingual sequence-to-sequence pretraining. *arXiv preprint arXiv:2108.01887*, 2021.
- Holger Schwenk. Filtering and mining parallel data in a joint multilingual space. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 228–234, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2037. URL <https://aclanthology.org/P18-2037>.
- Sotaro Takeshita, Tommaso Green, Niklas Friedrich, Kai Eckert, and Simone Paolo Ponzetto. X-scitldr: cross-lingual extreme summarization of scholarly documents. In *Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries*, pp. 1–12, 2022.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. Multilingual translation from denoising pre-training. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 3450–3466, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.304. URL <https://aclanthology.org/2021.findings-acl.304>.
- Sarubi Thillainathan, Surangika Ranathunga, and Sanath Jayasena. Fine-tuning self-supervised multilingual sequence-to-sequence models for extremely low-resource NMT. In *2021 Moratuwa Engineering Research Conference (MERCon)*, pp. 432–437, 2021. doi: 10.1109/MERCon52712.2021.9525720.
- Jörg Tiedemann and Santhosh Thottingal. OPUS-MT – building open translation services for the world. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pp. 479–480, Lisboa, Portugal, November 2020. European Association for Machine Translation. URL <https://aclanthology.org/2020.eamt-1.61>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Neha Verma, Kenton Murray, and Kevin Duh. Strategies for adapting multilingual pre-training for domain-specific machine translation. In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 31–44, 2022.
- Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu, and Michael Lyu. Understanding and improving sequence-to-sequence pretraining for neural machine translation. *arXiv preprint arXiv:2203.08442*, 2022.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin 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*, pp. 483–498, Online, 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.41. URL <https://aclanthology.org/2021.naacl-main.41>.

A EXPERIMENTS

A.1 DATASETS

Bible corpus Existing parallel corpora for Bible such as McCarthy et al. (2020), although multi-way parallel, have very little data for the languages we considered. Since we intend to perform a detailed analysis on dataset size, we curate a bible corpus for languages used in our experiments. We scrape Bible data from the web² and then automatically align the sentences (at verse level). We were able to curate a multi-way parallel corpus of size 25k for 4 languages (KN, GU, HI, TA). Note that Sinhala was scraped from a different website, thus has different content³.

Common Crawl (CC) CCAI (El-Kishky et al., 2020) corpus consists of parallel text that was automatically aligned using LASER sentence embeddings (Schwenk, 2018). The dataset, although noisy (Kreutzer et al., 2022), has been used to develop highly multilingual machine translation models such as M2M100 (Fan et al., 2020) and mBART multilingual MT (Tang et al., 2021).

PMIndia corpus (PMI) PMIndia (Haddow & Kirefu, 2020) is a parallel corpus for English and 13 other languages in India. It consists of news updates and excerpts of the Prime Minister’s speeches extracted from the Prime Minister of India’s website.

Government corpus (Gvt) The government document corpus (Fernando et al., 2020) is a multi-way parallel multilingual corpus for Sinhala, Tamil and English. It contains annual reports, committee reports, crawled content from government institutional websites, procurement documents, and acts from official Sri Lankan government documents.

FLORES The FLORES-101 (Goyal et al., 2021) dataset is a multilingual, multi-way parallel corpus whose sentences are extracted from English Wikipedia and translated into 101 languages. It consists of data from a variety of topics and domains. We use FLORESv1 (Guzmán et al., 2019) for Sinhala since it is not present in FLORES-101.

A.2 MODEL SETTINGS

We used HuggingFace libraries for our experiments. All experiments were run with seed 222 and performed using a Nvidia Volta of 32 GPU RAM.

We train up to 3 epochs with learning rate of $5 \cdot 10^{-5}$, dropout of 0.1, maximum length of 200 for the source and target, and a batch size of 10 for training and evaluation. We use the implementations in the HuggingFace Transformers library.

For kn language, we applied the related language fine-tuning strategy where we picked Telugu (te) for kn (Lee et al., 2022).

A.3 EVALUATION METRICS

We use SentencePiece BLEU (spBLEU in short), introduced by Goyal et al. (2022) as the evaluation metric for all our experiments. In this method, the BLEU scores are calculated for the text tokenized using sentence-piece subword model (which has been trained for all the 101 languages in FLORES-101 dataset). The standardization of tokenizers allows research to make comparisons among each

²Sinhala: <https://www.wordproject.org/bibles/si/index.htm>; and others: <https://ebible.org/download.php>

³We will be releasing the scripts to create the corpus on acceptance of the paper.

other. Further, Goyal et al. (2022) also show that spBLEU functions similar to BLEU and also has strong correlation with the tokenizer-independent Chrf++ metric (Popović, 2017). We use the official implementation provided in the sacreBLEU library⁴ (Post, 2018) for evaluating all the experiments.

A.4 JENSON-SHANNON DIVERGENCE

We calculate domain divergence using Jensen-Shannon (JS) divergence. JS divergence is calculated between two distributions P and Q using the formula:

$$JSD(P||Q) = \frac{1}{2}KL(P||M) + \frac{1}{2}KL(Q||M)$$

M represents an equally weighted sum of the two distributions (i.e., $M = \frac{1}{2}P + \frac{1}{2}Q$) and $KL(\cdot||\cdot)$ represents the Kullback–Leibler divergence.


The divergence is calculated for each (intermediate and fine-tuning) train and test set pair that were studied in the paper. The values are then averaged across all languages to obtain a similarity score that is independent of the target language. The train sets are composed of either 25k or 1k sentences and the test sets are composed entirely of 1k sentences. Each corpus is first tokenized using the NLTK package,⁵ stripped of all stopwords, and transformed into a (discrete) frequency distribution over all word tokens. We also convert all times and numbers into the tokens <TIME> and <NUMBER> respectively. The frequency distributions of each train and test set are then compared using the formula above. JS divergence ranges from 0 to 1 with lower values indicating that the two distributions are more similar.

⁴<https://github.com/mjpost/sacreBLEU>


⁵<https://www.nltk.org/>

B RESULTS

B.1 MBART VS MT5



(a) Difference in performance by training set



(b) Difference in performance by training set size

Figure 4: Comparative Analysis between mBART and mT5


B.2 BASELINE

		EN -> XX														
Training	Size	KN			GU			HI			SI			TA		
		FLORES	Bib	PMI	FLORES	Bib	PMI	FLORES	Bib	PMI	FLORES	Bib	Gvt	FLORES	Bib	Gvt
Zero shot	-	0.1	0.0	0.1	0.3	0.0	0.1	0.3	0.0	0.4	0.2	0.0	1.2	0.5	0.0	0.9
	1k	0.1	0.0	0.1	7.3	2.1	17.6	8.5	1.7	18.2	3.1	0.4	15.7	2.1	0.5	9.4
PMI/Gvt	10k	4.2	0.5	19.9	15.4	3.7	32.1	14.5	2.9	30.2	8.6	1.1	36.8	6.0	0.8	30.8
	25k	10.4	1.1	30.5	18.3	4.4	36.3	16.5	3.1	34.3	10.7	1.1	42.3	8.1	1.1	35.8
	50k	-	-	-	-	-	-	18.5	3.4	36.7	10.9	1.0	47.1	9.6	1.3	39.2
	100k	0.0	3.6	0.0	3.4	10.0	3.8	3.5	13.3	4.6	0.9	10.9	0.9	1.5	9.7	1.5
Bib	10k	0.6	17.3	0.4	4.6	23.2	3.5	3.8	26.5	3.1	1.6	30.9	1.2	2.2	24.0	1.6
	25k	1.3	24.9	0.5	4.5	27.3	3.2	3.2	30.5	2.4	1.7	36.1	1.0	2.1	29.4	1.1
	25k	0.2	0.0	0.1	7.8	1.1	4.0	15.1	5.6	13.0	9.3	2.1	14.6	9.4	5.7	7.8
CC	25k	0.2	0.1	2.1	7.5	0.9	6.4	23.1	6.8	19.8	15.1	4.0	25.9	16.0	8.4	15.2
	100k	2.8	0.1	2.1	7.5	0.9	6.4	23.1	6.8	19.8	15.1	4.0	25.9	16.0	8.4	15.2


		XX -> EN														
Training	Size	KN			GU			HI			SI			TA		
		FLORES	Bib	PMI	FLORES	Bib	PMI	FLORES	Bib	PMI	FLORES	Bib	Gvt	FLORES	Bib	Gvt
Zero Shot	-	0.1	0.0	0.0	0.3	0.0	0.2	0.3	0.0	0.5	0.2	1.2	0.0	0.3	0.0	0.5
	1k	0.4	0.0	2.4	9.6	2.8	17.2	12.7	4.4	20.5	6.2	1.9	16.2	5.4	1.7	13.2
PMI/Gvt	10k	1.5	0.5	12.3	13.7	2.8	17.2	16.6	5.2	32.0	7.3	2.7	30.7	8.4	2.5	29.5
	25k	2.3	0.7	14.9	13.8	4.0	31.8	18.4	5.0	36.5	9.0	2.5	37.1	8.2	2.5	34.3
	50k	-	-	-	-	-	-	19.3	5.4	38.6	9.2	2.6	41.0	10.0	3.0	39.4
	100k	0.0	0.8	0.0	4.0	15.6	3.3	6.8	18.1	6.6	2.9	15.0	41.0	3.5	13.2	2.3
Bib	10k	0.5	13.6	0.5	3.9	27.9	2.0	5.0	29.7	3.6	2.9	29.0	1.9	3.1	27.2	1.4
	25k	1.2	25.2	0.8	3.4	32.5	1.7	4.2	34.1	3.0	2.9	33.7	1.4	2.6	32.7	1.1
	25k	0.4	0.0	0.3	5.0	0.8	5.3	15.0	7.3	11.9	8.4	3.5	12.0	8.2	4.4	4.0
CC	25k	0.4	0.0	0.3	5.0	0.8	5.3	15.0	7.3	11.9	8.4	3.5	12.0	8.2	4.4	4.0
	100k	0.5	0.0	0.8	5.0	1.2	4.8	19.6	8.6	15.4	12.7	6.6	20.2	13.8	8.5	9.8

Table 5: Experimental results reported in spBLEU for the baseline. Zero shot corresponds to testing the mBART model without any finetuning.


B.3 INTERMEDIATE FINE-TUNING




(a) Intermediate Task Fine-Tuning with intermediate task - CC and final task - Bible. The spBLEU scores corresponds to scores on Bible test set.



(b) Intermediate Task Fine-Tuning with intermediate task - CC and final task - PMI/Gvt. The spBLEU scores corresponds to scores on PMI/Gvt test set.




(c) Intermediate Task Fine-Tuning with intermediate task - Bible and final task - PMI/Gvt. The spBLEU scores corresponds to scores on PMI/Gvt test set.




(d) Intermediate Task Fine-Tuning with intermediate task - PMI/Gvt and final task - Bible. The spBLEU scores corresponds to scores on Bible test set


Figure 5: Intermediate Task Fine-tuning - effect of dataset sizes used for fine-tuning for in-domain case. The Y-axis represents the spBLEU scores and X-axis shows the final task size. Colours show the intermediate task size (size 0 = baseline). Train Size 1 corresponds to size of the intermediate task data set and similarly Train Size 2 for final task.




(a) Intermediate Task Fine-Tuning with intermediate task - Bible and final task - PMI. The spBLEU scores corresponds to scores on FLORES test set.



(b) Intermediate Task Fine-Tuning with intermediate task - CC and final task - PMI. The spBLEU scores corresponds to scores on FLORES test set.



(c) Intermediate Task Fine-Tuning with intermediate task - PMI and final task - Bible. The spBLEU scores corresponds to scores on FLORES test set.



(d) Intermediate Task Fine-Tuning with intermediate task - CC and final task - Bible. The spBLEU scores corresponds to scores on FLORES test set.

Figure 6: Intermediate Task Fine-tuning - effect of dataset sizes used for fine-tuning for out-domain case. The Y-axis represents the spBLEU scores and X-axis shows the final task size. Colours show the intermediate task size (size 0 = baseline). Train Size 1 corresponds to size of the intermediate task data set and similarly Train Size 2 for final task.