

Enhancing Translation for Indigenous Languages: Experiments with Multilingual Models

Atnafu Lambebo Tonja¹, Hellina Hailu Nigatu², Olga Kolesnikova¹,
Grigori Sidorov¹, Alexander Gelbukh¹, Jugal Kalita³

¹Instituto Politécnico Nacional (IPN), Mexico,

²University of California, Berkeley, USA,

³University of Colorado, Colorado Springs, USA

Abstract

This paper describes CIC NLP’s submission to the AmericasNLP 2023 Shared Task on machine translation systems for indigenous languages of the Americas. We present the system descriptions for three methods. We used two multilingual models, namely M2M-100 and mBART50, and one bilingual (one-to-one) — Helsinki NLP Spanish-English translation model, and experimented with different transfer learning setups. We experimented with 11 languages from America and report the setups we used as well as the results we achieved. Overall, the mBART setup was able to improve upon the baseline for three out of the eleven languages.

1 Introduction

While machine translation systems have shown commendable performance in recent years, the performance is lagging for low-resource languages (Hadgu et al., 2022; Tonja et al., 2023). Since low-resource languages suffer from a lack of sufficient data (Siddhant et al., 2022; Haddow et al., 2022), most models and methods that are developed for high-resource languages do not work well in low-resource settings. Additionally, low-resource languages are linguistically diverse and have divergent properties from the mainstream languages in NLP studies (Zheng et al., 2021).

Though low-resource languages lack sufficient data to train large models, some such languages still have a large number of native speakers (Zheng et al., 2021). While the availability of language technologies such as machine translation systems can be helpful for such linguistic communities, they could also bring harm and exposure to exploitation (Hovy and Spruit, 2016). Borrowing from human-computer interaction (HCI) studies (Schneider et al., 2018), we want to acknowledge our belief that low-resource language speakers should be empowered to create technologies that benefit their communities. Many indigenous communi-

ties have community-rooted efforts for preserving their languages and building language technologies for their communities¹ and we hope that methods from Shared Tasks like this will contribute to their efforts.

Improving machine translation systems for low-resource languages is an active research area and different approaches (Zoph et al., 2016; Karakanta et al., 2018; Ortega et al., 2020a; Goyal et al., 2020; Tonja et al., 2022; Imankulova et al., 2017) have been to improve the performance of systems geared forward low-resource languages. We participated in the AmericasNLP 2023 Shared Task in hopes of contributing new approaches for low-resource machine translation that are likely to be helpful for community members interested in developing and adapting these technologies for their languages.

In recent years, large pre-trained models have been used for downstream NLP tasks, including machine translation (Brants et al., 2007) because of the higher performance in downstream tasks compared to traditional approaches (Han et al., 2021). One trend is to use these pre-trained models and fine-tune them on smaller data sets for specific tasks (Sun et al., 2019). This method has shown promising results in downstream NLP tasks for languages with low or limited resources (Tars et al., 2022; Zhao and Zhang, 2022). In our experiments, we used multilingual and bilingual models and employed different fine-tuning strategies for the eleven languages in the 2023 Shared Task (Ebrahimi et al., 2023).

In this paper, we describe the system setups we used and the results we obtained from our experiments. One of our systems improves upon the baseline for three languages. We also reflect on the setups we experimented with but ended up not submitting in hopes that future work could improve upon them.

¹<https://papareo.nz/>

2 Languages and Datasets

In this section, we present the languages and datasets used in our shared task submission. Table 1 provides an overview of the languages, their linguistic families, and the numbers of parallel sentences.

Language	ISO	Family	Train	Dev	Test
Aymara	aym	Aymaran	6,531	996	1,003
Bribri	bzd	Chibchan	7,508	996	1,003
Asháninka	cni	Arawak	3,883	883	1002
Chatino	czn	Zapotecan	357	499	1,000
Guarani	gn	Tupi-Guarani	26,032	995	1,003
Wixarika	hch	Uto-Aztecan	8,966	994	1003
Nahuatl	nah	Uto-Aztecan	410,000	672	996
Hñähñu	oto	Oto-Manguean	4,889	599	1,001
Quechua	quy	Quechuan	125,008	996	1,003
Shipibo-Konibo	shp	Panoan	14,592	996	1,002
Rarámuri	tar	Uto-Aztecan	14721	995	1002

Table 1: This table provides information about the languages with which we experimented including ISO language code and language family as well as the number of sentences in training, development, and test sets for each language.

Aymara is an Aymaran language spoken by the Aymara people of the Bolivian Andes. It is one of only a handful of Native American languages with over one million speakers (Homola, 2012). Aymara, along with Spanish and Quechua, is an official language in Bolivia and Peru. The data for the Aymara-Spanish come from the Global Voices (Tiedemann, 2012).

Bribri The Bribri language is spoken in Southern Costa Rica. Bribri has two major orthographies: Jara² and Constenla³ and the writing is not standardized which results in spelling variations across documents. In this case, the sentences use an intermediate representation to unify existing orthographies. The Bribri-Spanish data (Feldman, 2020) came from six different sources.

Asháninka Asháninka is an Arawakan language spoken by the Asháninka people of Peru and Acre, Brazil⁴. It is primarily spoken in the Satipo Province located in the Amazon forest. The parallel data for Asháninka-Spanish come mainly from three sources (Cushimariano Romano and Sebastián Q., 2008; Ortega et al., 2020b; Mihás, 2011) and translations by Richard Castro.

²<https://www.lenguabribri.com/se-tt%C3%B6-bribri-ie-hablemos-en-bribri>

³<https://editorial.ucr.ac.cr/index.php>

⁴<https://www.everyculture.com/wc/Norway-to-Russia/Ash-ninka.html>

Chatino Chatino is a group of indigenous Mesoamerican languages. These languages are a branch of the Zapotecan family within the Oto-Manguean language family. They are natively spoken by 45,000 Chatino people (Cruz and Woodbury, 2006) whose communities are located in the southern portion of the Mexican state of Oaxaca. The parallel data for Chatino-Spanish can be accessed here⁵.

Guarani Guarani is a South American language that belongs to the Tupi-Guarani family (Britton, 2005) of the Tupian languages. It is one of the official languages of Paraguay (along with Spanish), where it is spoken by the majority of the population, and where half of the rural population are monolingual speakers of the language (Mortimer, 2006).

Wixarika Wixarika is an indigenous language of Mexico that belongs to the Uto-Aztecan language family (de la Federación, 2003). It is spoken by the ethnic group widely known as the Huichol (self-designation Wixaritari), whose mountainous territory extends over portions of the Mexican states of Jalisco, San Luis Potosí, Nayarit, Zacatecas, and Durango, but mostly in Jalisco. United States: La Habra, California; Houston, Texas.

Nahuatl Nahuatl is a Uto-Aztecan language and was spoken by the Aztec and Toltec civilizations of Mexico⁶. The Nahuatl language has no standard orthography and has wide dialectical variations (Zheng et al., 2021).

Hñähñu Hñähñu, also known as Otomí, belongs to the Oto-Pamean family and lived in central Mexico for many centuries (Lastra, 2001). Otomí is a tonal language with a Subject-Verb-Object (SVO) word order (Ebrahimi et al., 2022). It is spoken in several states across Mexico.

Quechua The Quechua-Spanish data (Agić and Vulić, 2019; Tiedemann, 2012) has three different sources: the Jehova’s Witnesses texts, the Peru Minister of Education, and dictionary entries and samples collected by Diego Huarcaya. The Quechua language, also known as Runasimi is spoken in Peru and is the most widely spoken pre-Columbian language family of the Americas (Ebrahimi et al., 2022).

⁵<https://scholarworks.iu.edu/dspace/handle/2022/21028>

⁶www.elalliance.org/languages/nahuatl




Figure 1: Experiments on (a) fine-tuning bilingual model, (b) and (c) fine-tuning multilingual models. For (a) we fine-tuned the bilingual Spanish-English model on Spanish-Indigenous pairs, for (b) we fine-tuned the multilingual mBART50 model on Spanish-Indigenous pairs, and for (c) we fine-tuned the multilingual m2m100-48 model first on Spanish-All to produce m2m100-48_{inter} model and then fine-tuned the m2m100-48_{inter} model on Spanish-Indigenous pairs.

Shipibo-Konibo Shipibo-Konibo - Spanish data (Montoya et al., 2019; Galarreta et al., 2017) come from three different sources: samples from flashcards translated to Shipibo-Konibo, sentences translated from books for bilingual education, and dictionary entries.

Rarámuri Rarámuri, also known as Tarahumara is a Uto-Aztec language spoken in Northern Mexico (Caballero, 2017). Rarámuri is a polysynthetic and agglutinative language spoken mainly in the Sierra Madre Occidental region of Mexico (Ebrahimi et al., 2022).

3 Models

We experimented with two multilingual and one bilingual translation model with different transfer learning setups. We used M2M-100 and mBART50 for the multilingual experiment and the Helsinki-NLP Spanish-English model for the bilingual experiment. Figure 1 shows the models used in this experiment.

3.1 Bilingual models

For the bilingual model, as shown in Figure 1a, we use a publicly available Spanish - English⁷ pre-trained model from Huggingface⁸ trained by Helsinki-NLP. The pre-trained MT models released by Helsinki-NLP are trained on OPUS, an open-source parallel corpus for covering 500 languages (Tiedemann and Thottingal, 2020; Tiedemann, 2020). This model is trained using the framework of Marian NMT (Junczys-Dowmunt et al., 2018). Each model has six self-attention layers in the encoder and decoder parts, and each layer has eight attention heads.

⁷<https://huggingface.co/Helsinki-NLP/opus-mt-es-en>

⁸<https://huggingface.co/>

We used this model with the intention that the model trained with high-resource languages will improve the translation performance of low-resource indigenous languages when using a model trained with high-resource languages. We fine-tuned the Spanish-English model for each of the Spanish-to-Indigenous language pairs.

3.2 Multilingual models

For multilingual models, we used the Many-to-Many multilingual translation model that can translate directly between any pair of 100 languages (M2M100) (Fan et al., 2021) with 48M parameters and a sequence-to-sequence denoising auto-encoder pre-trained on large-scale monolingual corpora in 50 languages (mBART50) (Tang et al., 2020). We fine-tuned multilingual models in two ways:

1. We fine-tuned two multilingual models on each Spanish-Indigenous language pair for 5 epochs and evaluated their performance using the development data before training the final submission system. As shown in Figure 1b, for the final system, we only fine-tuned mBART50 on Spanish-indigenous data based on the development set evaluation performance.
2. Fine-tuning multilingual models first on the Spanish - All (mixture of all indigenous language data) dataset to produce an intermediate model and then fine-tuning the intermediate model for each of the Spanish-Indigenous language pairs as shown in Figure 1c. For this experiment, we combined all language pairs' training data to form a Spanish - all parallel corpus, and then we first fine-tuned m2m100-48 using a combined dataset for five

Data	Model	aym	bzd	cni	czn	gn	hch	nah	oto	quy	shp	tar	Average
Baseline		28.3	16.5	25.8	-	33.6	30.4	26.6	14.7	34.3	32.9	18.4	-
Dev	M1	12.25	20.3	26.65	-	23.83	11.09	29.55	6.57	35.04	20.99	14.12	20.03
	M2	20.65	18.59	20.63	-	20.40	12.7	18.66	10.17	33.53	21.03	13.54	18.99
	M3	14.70	19.9	25.62	-	23.62	11.82	29.94	7.94	35.3	21.32	14.19	20.43
	M4	20.89	12.17	23.59	-	20.84	13.51	22.63	7.16	30.86	18.02	12.60	18.22
Test	M2	19.05	19.90	23.50	14.41	19.35	12.05	21.88	9.22	34.15	20.43	13.86	18.89
	M3	18.52	21.17	25.85	15.61	21.75	13.88	26.57	7.40	35.62	21.26	14.87	20.22
	M4	18.59	13.24	23.79	13.64	20.94	14.67	22.60	7.28	32.75	18.13	12.07	17.97

Table 2: chrF2 scores for the three submissions, computed on the development and test sets. M1, M2, M3, and M4 represent M2M100-48, M2M100-48_{inter}, mBART50 and Helsinki-NLP models respectively. The development set evaluations are used to select the best-performing model before working on submission data. The development set was not trained when evaluating the dev set, but we included the dev set during training for the final submission. The **bold** results show the models that out-performs the baseline (Vázquez et al., 2021) results. The **bold** results show out-performing models from our three model setups(excluding the baseline) for each individual language.

epochs and saved the model, here referred to as m2m100-48_{inter} model. We fine-tuned the m2m100-48_{inter} model again on each Spanish-Indigenous language pair for another 5 epochs and evaluated the performance on the development set before training the final submission system.

Evaluation We used chrF2 (Popović, 2017) evaluation metric to evaluate our MT systems.

4 Results

We submitted three (two multilingual and one bilingual) systems, as shown in Table 2, namely m2m100-48_{inter}, mBART50, and Helsinki-NLP. We included the dev set performance for all the models we trained before the final model to compare the results with the final model evaluated by using test set data. From the dev set result, it can be seen that fine-tuning the multilingual model on the Spanish-Indigenous language pair outperforms the fine-tuned result of the bilingual and m2m100-48_{inter} models. From all the models evaluated using the dev set, mBART50 outperformed the others on average.

Our test results show comparable results when compared to the strongest baseline shared by the AmericasNLP 2023, and our model outperformed the baseline for Spanish-Bribri (es-bzd), Spanish-Asháninka (es-cni), and Spanish-Quechua (es-quy) pairs. Similarly, mBART50 outperformed the other models on average on the test set.

5 Conclusion

In this work, we present the system descriptions and results for our submission to the 2023 AmericasNLP Shared Task on Machine Translation into

Indigenous Languages. We used pre-trained models and tested different fine-tuning strategies for the eleven languages provided for the shared task. We used one bilingual (Helsinki NLP English-Spanish model) and two multilingual (M2M-100 and mBART50) models for our experiments. In addition to fine-tuning the individual languages’ data, we concatenated the data from all eleven languages to create a Spanish-All dataset and fine-tuned the M2M-100 model before fine-tuning for the individual languages. Our mBAERT50 model beat the strong baseline in three languages.

Acknowledgements

References

- Željko Agić and Ivan Vulić. 2019. [JW300: A wide-coverage parallel corpus for low-resource languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210, Florence, Italy. Association for Computational Linguistics.
- Thorsten Brants, Ashok C. Popat, Peng Xu, Franz J. Och, and Jeffrey Dean. 2007. [Large language models in machine translation](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 858–867, Prague, Czech Republic. Association for Computational Linguistics.
- A Scott Britton. 2005. Guaraní-english/english-guaraní: Concise dictionary. *New York: Hippocrene*.
- Gabriela Caballero. 2017. [Choguita rarámuri \(tarahumara\) language description and documentation: a guide to the deposited collection and associated materials. language documentation conservation](#).
- Emiliana Cruz and Anthony C Woodbury. 2006. The sandhi of tones in the chatino de quiahije. In *Memoirs*

- of the Congress of Indigenous Languages of Latin America-III*.
- Rubén Cushimariano Romano and Richer C. Sebastián Q. 2008. Ñaantsipeta asháninkaki birakochaki. diccionario asháninka-castellano. versión preliminar. <http://www.lengamer.org/publicaciones/diccionarios/>.
- Diario Oficial de la Federación. 2003. Ley general de derechos lingüísticos de los pueblos indígenas. *Dirección General de Bibliotecas*.
- Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Meza-Ruiz, Gustavo A. Giménez-Lugo, Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Ngoc Thang Vu, and Katharina Kann. 2022. *Americasnli: Evaluating zero-shot natural language understanding of pretrained multilingual models in truly low-resource languages*.
- Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Enora Rice, Cynthia Montañó, John Ortega, Shruti Rijhwani, Alexis Palmer, Rolando Coto-Solano, Hilaria Cruz, and Katharina Kann. 2023. Findings of the AmericasNLP 2023 shared task on machine translation into indigenous languages. In *Proceedings of the Third Workshop on Natural Language Processing for Indigenous Languages of the Americas*. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. Beyond english-centric multilingual machine translation. *The Journal of Machine Learning Research*, 22(1):4839–4886.
- Coto-Solano R Feldman, I. 2020. *Neural Machine Translation Models with Back-Translation for the Extremely Low-Resource Indigenous Language Bribri*. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3965–3976.
- Ana-Paula Galarreta, Andrés Melgar, and Arturo Oncevay. 2017. *Corpus creation and initial SMT experiments between Spanish and Shipibo-konibo*. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 238–244, Varna, Bulgaria. INCOMA Ltd.
- Vikrant Goyal, Sourav Kumar, and Dipti Misra Sharma. 2020. Efficient neural machine translation for low-resource languages via exploiting related languages. In *Proceedings of the 58th annual meeting of the association for computational linguistics: student research workshop*, pages 162–168.
- Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. Survey of Low-Resource Machine Translation. *Computational Linguistics*, 48(3):673–732.
- Asmelash Teka Hadgu, Abel Aregawi, and Adam Beaudoin. 2022. *Lesan – machine translation for low resource languages*. In *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pages 297–301. PMLR.
- Xu Han, Zhengyan Zhang, Ning Ding, Yuxian Gu, Xiao Liu, Yuqi Huo, Jiezhong Qiu, Yuan Yao, Ao Zhang, Liang Zhang, et al. 2021. Pre-trained models: Past, present and future. *AI Open*, 2:225–250.
- Petr Homola. 2012. Building a formal grammar for a polysynthetic language. In *Formal Grammar: 15th and 16th International Conferences, FG 2010, Copenhagen, Denmark, August 2010, FG 2011, Ljubljana, Slovenia, August 2011, Revised Selected Papers*, pages 228–242. Springer.
- Dirk Hovy and Shannon L. Spruit. 2016. *The social impact of natural language processing*. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- Aizhan Imankulova, Takayuki Sato, and Mamoru Komachi. 2017. Improving low-resource neural machine translation with filtered pseudo-parallel corpus. In *Proceedings of the 4th Workshop on Asian Translation (WAT2017)*, pages 70–78.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. *Marian: Fast neural machine translation in C++*. In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Alina Karakanta, Jon Dehdari, and Josef van Genabith. 2018. Neural machine translation for low-resource languages without parallel corpora. *Machine Translation*, 32:167–189.
- Y. Lastra. 2001. *Chapter 6. Otomí Language Shift and Some Recent Efforts to Reverse It*, pages 142–165. Multilingual Matters, Bristol, Blue Ridge Summit.
- Elena Mihás. 2011. *Añaani katonkosatzi parenini, El idioma del alto Perené*. Milwaukee, WI: Clarks Graphics.
- Héctor Erasmo Gómez Montoya, Kervy Dante Rivas Rojas, and Arturo Oncevay. 2019. *A continuous improvement framework of machine translation for Shipibo-konibo*. In *Proceedings of the 2nd Workshop on Technologies for MT of Low Resource Languages*, pages 17–23, Dublin, Ireland. European Association for Machine Translation.
- Katherine Mortimer. 2006. Guaraní académico or jopará? educator perspectives and ideological debate in paraguayan bilingual education. *Working Papers in Educational Linguistics (WPEL)*, 21(2):3.

- John E Ortega, Richard Castro Mamani, and Kyunghyun Cho. 2020a. Neural machine translation with a polysynthetic low resource language. *Machine Translation*, 34(4):325–346.
- John E Ortega, Richard Alexander Castro-Mamani, and Jaime Rafael Montoya Samame. 2020b. [Overcoming resistance: The normalization of an Amazonian tribal language](#). In *Proceedings of the 3rd Workshop on Technologies for MT of Low Resource Languages*, pages 1–13, Suzhou, China. Association for Computational Linguistics.
- Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the second conference on machine translation*, pages 612–618.
- Hanna Schneider, Malin Eiband, Daniel Ullrich, and Andreas Butz. 2018. [Empowerment in hci - a survey and framework](#). CHI '18, page 1–14, New York, NY, USA. Association for Computing Machinery.
- Aditya Siddhant, Ankur Bapna, Orhan Firat, Yuan Cao, Mia Xu Chen, Isaac Caswell, and Xavier Garcia. 2022. [Towards the next 1000 languages in multilingual machine translation: Exploring the synergy between supervised and self-supervised learning](#).
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*, pages 194–206. Springer.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Maali Tars, Taïdo Purason, and Andre Tättar. 2022. Teaching unseen low-resource languages to large translation models. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 375–380.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Jörg Tiedemann. 2020. [The tatoeba translation challenge – realistic data sets for low resource and multilingual MT](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1174–1182, Online. Association for Computational Linguistics.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Atnafu Lambebo Tonja, Olga Kolesnikova, Muhammad Arif, Alexander Gelbukh, and Grigori Sidorov. 2022. Improving neural machine translation for low resource languages using mixed training: The case of ethiopian languages. In *Advances in Computational Intelligence: 21st Mexican International Conference on Artificial Intelligence, MICAI 2022, Monterrey, Mexico, October 24–29, 2022, Proceedings, Part II*, pages 30–40. Springer.
- Atnafu Lambebo Tonja, Olga Kolesnikova, Alexander Gelbukh, and Grigori Sidorov. 2023. [Low-resource neural machine translation improvement using source-side monolingual data](#). *Applied Sciences*, 13(2).
- Raúl Vázquez, Yves Scherrer, Sami Virpioja, and Jörg Tiedemann. 2021. The helsinki submission to the americasnlp shared task. In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*. The Association for Computational Linguistics.
- Jing Zhao and Wei-Qiang Zhang. 2022. Improving automatic speech recognition performance for low-resource languages with self-supervised models. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1227–1241.
- Francis Zheng, Machel Reid, Edison Marrese-Taylor, and Yutaka Matsuo. 2021. Low-resource machine translation using cross-lingual language model pre-training. In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 234–240.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. *arXiv preprint arXiv:1604.02201*.