

IndT5: A Text-to-Text Transformer for 10 Indigenous Languages

El Moatez Billah Nagoudi¹, Wei-Rui Chen¹, Muhammad Abdul-Mageed¹, Hasan Cavusoglu²

¹ Natural Language Processing Lab,

^{1,2} The University of British Columbia

¹ {moatez.nagoudi, weirui.chen, muhammad.mageed}@ubc.ca, ² cavusoglu@sauder.ubc.ca

Abstract

Transformer language models have become fundamental components of natural language processing based pipelines. Although several Transformer models have been introduced to serve many languages, there is a shortage of models pre-trained for low-resource and Indigenous languages. In this work, we introduce IndT5, the first Transformer language model for Indigenous languages. To train IndT5, we build IndCorpus—a new dataset for ten Indigenous languages and Spanish. We also present the application of IndT5 to machine translation by investigating different approaches to translate between Spanish and the Indigenous languages as part of our contribution to the AmericasNLP 2021 Shared Task on Open Machine Translation. IndT5 and IndCorpus are publicly available for research.¹

1 Introduction

Indigenous languages are starting to attract attention in the field of natural language processing (NLP), with the number of related publications growing in recent years (Mager et al., 2018). In spite of this interest, there remains a multitude of challenges for handling Indigenous languages. Complexity of the morphological systems of some of these languages and lack of standard orthography for writing them are among these challenges (Mager et al., 2018; Littell et al., 2018). The most fundamental issue facing NLP efforts, however, remains the lack of digital textual data that can be exploited for systems development.

In this work, we describe a scenario usually faced when trying to develop NLP systems for Indigenous languages and we focus on machine translation (MT). We adopt a neural machine translation approach (NMT) (Koehn, 2017) as our method. We show that, in spite of its recent success on many

¹<https://github.com/UBC-NLP/IndT5>

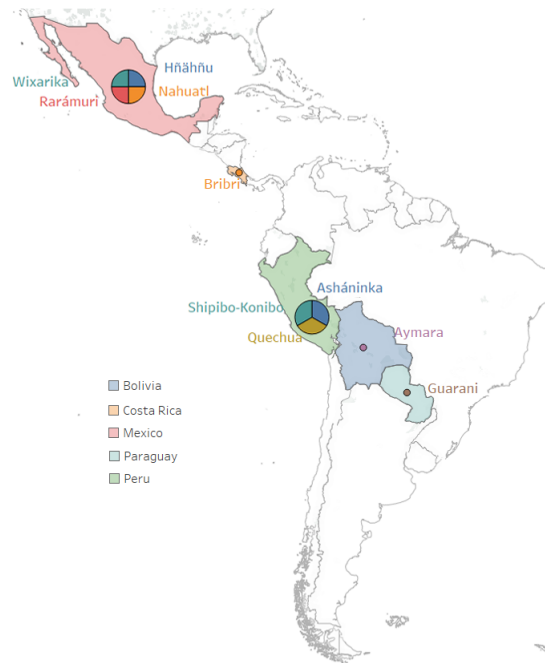


Figure 1: A map of the ten Indigenous languages covered by IndT5, our text-to-text Transformer model, and our IndCorpus dataset. The languages are mainly spoken in five Latin American countries.

contexts, NMT still struggles in very low-resource settings involving Indigenous languages. This is due to the core difficulty of lack of parallel textual data, but also even monolingual data.

Although our main goal in this work in particular is to develop translation models from Spanish to several Indigenous languages of the Americas, we adopt a transfer learning approach where we offer resources that can be exploited for other downstream tasks. Namely, we build a dataset for ten Indigenous languages and Spanish which we refer to as **IndCorpus**. Figure 1 and Table 1 provide an overview of the ten Indigenous languages in our new dataset (Eberhard et al., 2021). We also exploit **IndCorpus** for pre-training a Transformer language model following the unified approach introduced by (Raffel et al., 2019). Our resulting model,

Language	Code	Main location	Speakers
Aymara	aym	Bolivia	1,677,100
Asháninka	cni	Peru	35,200
Bribri	bzd	Costa Rica	7,000
Guarani	gn	Paraguay	6,652,790
Hñähñu	oto	Mexico	88,500
Nahuatl	nah	Mexico	410,000
Quechua	quy	Peru	7,384,920
Rarámuri	tar	Mexico	9,230
Shipibo-Konibo	shp	Peru	22,500
Wixarika	hch	Mexico	52,500

Table 1: Overview of our ten Indigenous languages (Eberhard et al., 2021).

IndT5, treats every text NLP problem as a “text-to-text” problem, i.e. taking text as input and producing new text as output. We apply **IndT5** to the MT task as a way to transfer knowledge acquired by the model to this particular context. Our experiments show the utility of our new language model and the dataset it exploits for the downstream Indigenous MT task but that very large space for improvement still exists.

The rest of the paper is organized as follows: In Section 2, we introduce recent MT work in low-resource and Indigenous languages settings. In Section 3, we describe how we develop our new language model for ten Indigenous languages. In Section 4, we describe our NMT models. We conclude in Section 5.

2 Related Work

2.1 Low-Resource MT

A number of methods and techniques have been proposed to mitigate the effects of having rather small datasets for machine translation. These include data augmentation, transfer learning, hyperparameter tuning, incorporating linguistic knowledge, and knowledge distillation.

Since the main bottleneck of low-resource MT is the lack of abundant parallel textual data, data augmentation is straightforwardly a potential method to enhance the model performance. Back translation is a way to augment parallel data (Sennrich et al., 2016a). By training a target-to-source translation model with original data and feeding in monolingual data of target language, synthetic parallel data is generated. If the target language is rich in textual data, much synthetic parallel data can be added into training data and may benefit the final translation model.

Transfer learning is another method that can boost the performance of MT on low-resource languages (Zoph et al., 2016; Nguyen and Chiang, 2017; Kocmi and Bojar, 2018). The rationale behind one approach to transfer learning is that knowledge obtained while translating high-resource languages may be transferable to translation of low-resource languages. In Zoph et al. (2016), a parent model is first trained on a high-resource language pair (i.e., French to English) then a child model is trained on a low-resource language pair (i.e., Uzbek to English). The Uzbek-English model has 10.7 BLEU score without parent model and 15.0 with the parent model. It is also shown that the more similar the two source languages, the more performance gain is possible. For example, a Spanish-English MT model has 16.4 BLEU score without parent model and 31.0 with French-English parent model. The performance gain is much more than when transferring French-English parent model to the more distant context of the Uzbek-English child model.

Sennrich and Zhang (2019) argue that instead of using hyperparameters that work in high-resource settings, there should be a set of hyperparameters specific to the low-resource scenario. For example, keeping the vocabulary size small, training a model with relatively small capacity, and having smaller batch size may be beneficial to model performance. When building a vocabulary with BPE, by reducing the the number of merge operations, a smaller vocabulary can be obtained and an inclusion of low-frequency (sub)words can be avoided. Inclusion of inclusion of low-frequency (sub)words could otherwise negatively influencing representation learning effectiveness.

Leveraging linguistic knowledge for data augmentation, Zhou et al. (2019) use a rule-based syntax parser and a dictionary to generate parallel data. By reordering target-language sentences into source-language syntactic structure and then mapping target-language words into source-language words with a dictionary, the size of parallel data is enlarged and translation performance is improved.

Baziotis et al. (2020) leverage a language model to help enhance the performance of the translation model. Similar to the idea of knowledge distillation (Hinton et al., 2015), a teacher model and a student model are trained where the language model plays the role of teacher and translation model plays the role of student. With this design, the teacher model

needs only monolingual data and does not have to rely on large parallel data.

2.2 MT of Indigenous Languages

Unlike high-resource languages such as English and French, Indigenous languages are often low-resource. Due to this, it is common that researchers of Indigenous languages adopt methods that can fare well in low-resource scenarios. This includes using the Transformer architecture and its variants in both low-resource (Adebara et al., 2021, 2020; Przystupa and Abdul-Mageed, 2019) and Indigenous language (Feldman and Coto-Solano, 2020; Orife, 2020; Le and Sadat, 2020) settings.

Despite the fact that Indigenous languages face difficulties similar to most low-resource languages, there are some challenges specific to Indigenous languages. As Mager et al. (2018) point out, some Indigenous languages have complex morphological systems and some have various non-standardized orthographic conventions. For example, Micher (2018) shows that in Inuktitut, an Indigenous language in North America with a complex morphological system, a corpus of one million tokens, there are about 225K different types for Inuktitut while about 30K types for English. Also, Micher (2018) shows that there can be lack of standardized spelling for some words. For example, the word *Haammalat* in Inuktitut has another seven different forms.

To cope with the issue of complex morphology, Ortega et al. (2020) build a translation model for Quechua, an Indigenous language of South America, with an integrated morphological segmentation method. To treat orthographic variation, Feldman and Coto-Solano (2020) standardize text with a rule-based system which converts diacritics and letters to contemporary orthographic convention.

3 IndT5

We train an Indigenous language model adopting the unified and flexible text-to-text transfer Transformer (T5) approach (Raffel et al., 2019). T5 treats every text-based language task as a “text-to-text” problem, taking text format as input and producing new text format as output. T5 is essentially an encoder-decoder Transformer (Vaswani et al., 2017), with the encoder and decoder similar in configuration and size to a BERT_{Base} (Devlin et al., 2019) but with some architectural modifica-

tions. Modifications include applying a normalization layer before a sub-block and adding a pre-norm (i.e., initial input to the sub-block output). We call our resulting model **IndT5**. We now describe our dataset, vocabulary, and pre-training method for developing **IndT5**.

3.1 Training Data

We build **IndCorpus**, a collection of ten Indigenous languages and Spanish comprising 1.17 GB of text (~5.37M sentences), to pre-train **IndT5**. **IndCorpus** is collected from both Wikipedia and the Bible. Table 2 provides the size and number of sentences for each language in our dataset.

3.2 IndT5 Vocabulary

The T5 (Raffel et al., 2019) model is based on a vocabulary acquired by the SentencePiece library² using English, French, German, and Romanian web pages from “Colossal Clean Crawled Corpus” (or C4 for short). We use a similar procedure to create our Indigenous languages vocabulary. Namely, we use SentencePiece (Kudo, 2018) to encode text as WordPiece (Sennrich et al., 2016b) tokens with a vocabulary size of 100K WordPieces extracted from **IndCorpus**.

3.3 Unsupervised Pre-Training

We leverage our unlabeled Indigenous corpus, **IndCorpus**, to pre-train **IndT5**. For that, we use a denoising objective (Raffel et al., 2019) that does not require labels. The main idea is feeding the model with corrupted (masked) versions of the original sentence, and training it to reconstruct the original sentence. Inspired by BERT’s objective (i.e., masked language model) (Devlin et al., 2019), the denoising objective (Raffel et al., 2019) works by randomly sampling and dropping out 15% of tokens in the input sequence. All consecutive spans of dropped-out tokens are then replaced by a single sentinel token. We pre-train our model for 100K steps on the **IndCorpus** using the T5_{Base} architecture.³ We refer to this model as IndT5_{100k}. Afterwards, we further pre-train on only the ten Indigenous languages part of our dataset (i.e., without the Spanish data) for 40K steps. We refer to this version of the model as IndT5_{140k}. For both pre-training steps, we use a learning rate of 0.01,

²<https://github.com/google/sentencepiece>

³Both encoder and decoder of T5_{Base} model has 12 layers each with 12 attention heads, and 768 hidden units.

Target language	Wikipedia		Bible	
	Size (MB)	Sentences	Size (MB)	Sentences
Hñähñu	-	-	1.4	7.5K
Wixarika	-	-	1.3	7.5K
Nahuatl	5.8	61.1K	1.5	7.5K
Guarani	3.7	28.2K	1.3	7.5K
Bribri	-	-	1.5	7.5K
Rarámuri	-	-	1.9	7.5K
Quechua	5.9	97.3K	4.9	31.1K
Aymara	1.7	32.9K	5	30.7K
Shipibo-Konibo	-	-	1	7.9K
Asháninka	-	-	1.4	7.8K
Spanish	1.13K	5M	-	-
Total	1.15K	5.22M	19.8	125.3K

Table 2: Datasets in IndCorpus by language

Languages	Train	Dev	Test
es-aym	6, 531	996	1, 003
es-cni	3, 883	883	1, 003
es-bzd	7, 506	996	1, 003
es-gn	26, 032	995	1, 003
es-oto	4, 889	599	1, 003
es-nah	16, 145	672	1, 003
es-quy	125, 008	996	1, 003
es-tar	14, 720	995	1, 003
es-shp	14, 592	996	1, 003
es-hch	8, 966	994	1, 003

Table 3: Distribution of MT data

a batch size of 128 sequences, and a maximum sequence length of 512. We use the original implementation of T5 in the TensorFlow framework.⁴ We train the models on Google Cloud TPU with 8 cores (v3.8) from TensorFlow Research Cloud (TFRC).⁵

4 Our Machine Translation Models

4.1 Parallel Data

As part of the AmericasNLP 2021 Shared Task on Open Machine Translation, the training (Train) and development (Dev) datasets for ten target Indigenous languages along with the source language Spanish were released. All the datasets are manually translated. Table 3 shows the number of sentences of different language pairs in shared task

⁴<https://github.com/google-research/text-to-text-transfer-transformer>

⁵<https://www.tensorflow.org/tfrc>

data. Table 4 provides example sentences extracted from the Dev dataset with their corresponding translations.

4.2 Approach

For all languages pairs except *quy* and *gn*, we fine-tune each of the two versions of our language model, i.e., both IndT5_{100k} and IndT5_{140k}, under two conditions: **(A)** we train on Train using 100% of Dev data for validation, for 150 epochs; **(B)** we fine-tune the best epoch from setting A for 50 epochs, adding 80% of Dev data to Train (using the remaining 20% Dev for validation).

4.3 Evaluation

We report the results of both IndT5_{100k} and IndT5_{140k} models using two metrics: BLEU score (Papineni et al., 2002) and ChrF++ (Popović, 2017). Tables 5 and 6 show the results of both models on Test sets for each of the language pairs using settings A and B described in Section 4.2, respectively.

4.4 Discussion

The results presented in Table 5 and Table 6 show that all our models, with both settings A and B, outperform the respective baselines across all languages. An exception is the languages *aym* and *shp*. As expected, fine-tuning the IndT5_{100k} and IndT5_{140k} models using the training data and 80% of the Dev data (i.e., setting B) improves the results with a mean of +0.003% and +0.04% in ChrF++ on the Test data, respectively. Interestingly, fur-

Pair	Sentence	Translation
es-aym	Algunos actores usan el teatro comunitario para mejorar.	Yaqhip akturanakax juk' amp yatsuñatakiw ayllunkir tiyatruv mantapxi.
	Los artistas de IRT ayudan a los niños en las escuelas.	IRT artistanakax jisk'a yatiqañ utankir wawanakaruw yanapapxi.
es-cni	Pensé que habías ido al campamento.	Nokenkeshireashitaka pijaiti imabeyetinta.
	Viajar es un beneficio que obtenemos.	Akenayeeterika aparo ayeeti aneakeri.
es-bzd	Fui a un seminario que se hizo vía satélite.	Ye' dë'rö seminario ã wéx yö' satélite kī.
	El grupo está interesado en temas ambientales.	E' wakpa kī ujte kiänã e' dör káx ajkóqnük.
es-gn	Veía a su hermana todos los días.	Ko'êko'êre ohecha heindýpe.
	Ramona nunca ha estado en Concord.	Ramona noíriva Concord-pe.
es-nah	Santo trabajó para Disney y operó las tazas de té.	zanto quitequitiilih Disney huan quinpeontih in cafen caxitl
	La hermana de la abuela no era blanca.	ihueltiuh in cihtli ixchipahuac catca
es-quy	De vez en cuando me gusta comer ensalada.	Yananpiqa ensaladatum mikuytam munani
	Ellos vivían en Broad Street.	Broad Streetpi paykuna yacharqaku.
es-tar	Es un hombre griego.	Bilé rejói Griego ju
	Nuestro padre dijo que no los llamaran animales.	Kini onó aniyé mapu ke chuwé namúti anéba ajaré jákami.
es-shp	El Museo se ve afectado por las inversiones.	Ja Museora en oinai inversionesbaon afectana.
	Loren Field es el científico principal de la escuela	Nato Loren Field riki científico rekena axeti xobonko
es-hch	Era una selva tropical.	pe h+k+t+kai metsi+ra+ ye tsie nieka ti+x+kat+.
	Son más económicos porque son realmente buenos en gas.	p+ h+k+ nip+ka raye at+ka aix+ m+ anenek+ ik+ gas.

Table 4: Example sentences of the various language pairs and corresponding translations (from Dev set).

Pair	Baseline		Setting A		Setting B	
	Bleu	ChrF++	Bleu	ChrF++	Bleu	ChrF++
aym	0.3	0.188	1.01	0.178	0.76	0.186
cni	0.03	0.104	0.09	0.176	0.09	0.178
bzd	0.54	0.077	0.86	0.11	0.89	0.111
oto	0.01	0.059	0.03	0.081	0.04	0.083
nah	0.33	0.182	-	-	0.16	0.196
tar	0.01	0.046	0.06	0.102	-	-
hch	3.18	0.126	4.95	0.186	5.09	0.186

Table 5: Evaluation results of IndT5_{100k} in BLEU and ChrF++ on the Test sets for the different language pairs.

ther pre-training IndT5 on only the ten Indigenous languages (i.e. target languages) produces better results with an average improvement of +0.003% and +0.004% in settings A and B, respectively. Overall, the impact of limited data is clear.

5 Conclusion

In this work, we introduced a new Transformer language model (**IndT5**) and a dataset (**IndCorpus**) for ten Indigenous languages and Spanish. We applied **IndT5** to the MT task on eight languages pairs as part of our submission to the AmericasNLP 2021 Shared Task. While **IndT5** helps improve translation, the task remains hard due to absence of parallel as well as mono-

Pair	Baseline		Setting A		Setting B	
	Bleu	ChrF++	Bleu	ChrF++	Bleu	ChrF++
aym	0.3	0.188	0.820	0.182	0.990	0.190
cni	0.03	0.104	0.070	0.178	0.080	0.183
bzd	0.54	0.077	0.990	0.112	0.940	0.113
oto	0.01	0.059	0.030	0.082	0.040	0.084
nah	0.33	0.182	0.150	0.188	0.160	0.196
tar	0.01	0.046	0.080	0.102	0.050	0.105
shp	0.34	0.139	0.160	0.124	0.230	0.124
hch	3.18	0.126	5.100	0.194	5.520	0.195

Table 6: Evaluation results of IndT5_{140k} in BLEU and ChrF++ on the Test sets for the different language pairs.

lingual data. In the future, we plan to integrate statistical MT methods to augment our data as well as investigate best hyperparameters for our neural models.

Acknowledgements

We gratefully acknowledge support from the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada, Canadian Foundation for Innovation, Compute Canada, and UBC ARC-Sockeye. We also thank the Google TFRC program for providing us with free TPU access.

References

- Ife Adebara, Muhammad Abdul-Mageed, and Miikka Silfverberg. 2021. Translating the Unseen? Yorùbá-English MT in Low-Resource, Morphologically-Unmarked Settings. *AfricNLP*.
- Ife Adebara, El Moatez Billah Nagoudi, and Muhammad Abdul Mageed. 2020. [Translating similar languages: Role of mutual intelligibility in multilingual transformers](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 381–386, Online. Association for Computational Linguistics.
- Christos Baziotis, Barry Haddow, and Alexandra Birch. 2020. Language model prior for low-resource neural machine translation. *arXiv preprint arXiv:2004.14928*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2021. [Ethnologue: Languages of the world, twenty-fourth edition](#). Dallas, Texas. SIL International.
- Isaac Feldman and Rolando Coto-Solano. 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.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Tom Kocmi and Ondrej Bojar. 2018. [Trivial transfer learning for low-resource neural machine translation](#). *CoRR*, abs/1809.00357.
- Philipp Koehn. 2017. Neural machine translation. *arXiv preprint arXiv:1709.07809*.
- Taku Kudo. 2018. Subword regularization: Improving neural network translation models with multiple subword candidates. *arXiv preprint arXiv:1804.10959*.
- N. Tan Le and F. Sadat. 2020. Addressing challenges of indigenous languages through neural machine translation: The case of inuktitut-english.
- Patrick Littell, Anna Kazantseva, Roland Kuhn, Aidan Pine, Antti Arppe, Christopher Cox, and Marie-Odile Junker. 2018. [Indigenous language technologies in Canada: Assessment, challenges, and successes](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2620–2632, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Manuel Mager, Ximena Gutierrez-Vasques, Gerardo Sierra, and Ivan Meza-Ruiz. 2018. [Challenges of language technologies for the indigenous languages of the Americas](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 55–69, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- J.C. Micher. 2018. *Addressing Challenges of Machine Translation of Inuit Language*. ARL-TN. US Army Research Laboratory.
- Toan Q. Nguyen and David Chiang. 2017. [Transfer learning across low-resource, related languages for neural machine translation](#). *CoRR*, abs/1708.09803.
- Iroro Orife. 2020. Towards neural machine translation for edoid languages. *arXiv preprint arXiv:2003.10704*.
- John E. Ortega, Richard Castro Mamani, and Kyunghyun Cho. 2020. [Neural machine translation with a polysynthetic low resource language](#). *Machine Translation*, 34(4):325–346.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the second conference on machine translation*, pages 612–618.
- Michael Przystupa and Muhammad Abdul-Mageed. 2019. Neural machine translation of low-resource and similar languages with backtranslation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 224–235.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

- Rico Sennrich and Biao Zhang. 2019. Revisiting low-resource neural machine translation: A case study. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010.
- Chunting Zhou, Xuezhe Ma, Junjie Hu, and Graham Neubig. 2019. [Handling syntactic divergence in low-resource machine translation](#).
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575.