

Multilingual Neural Machine Translation: Can Linguistic Hierarchies Help?

Fahimeh Saleh Wray Buntine Gholamreza Haffari Lan Du*

first.last@monash.edu

Monash University

Abstract

Multilingual Neural Machine Translation (MNMT) trains a single NMT model that supports translation between multiple languages, rather than training separate models for different languages. Learning a single model can enhance the low-resource translation by leveraging data from multiple languages. However, the performance of an MNMT model is highly dependent on the type of languages used in training, as transferring knowledge from a diverse set of languages degrades the translation performance due to negative transfer. In this paper, we propose a Hierarchical Knowledge Distillation (HKD) approach for MNMT which capitalises on language groups generated according to typological features and phylogeny of languages to overcome the issue of negative transfer. HKD generates a set of multilingual teacher-assistant models via a selective knowledge distillation mechanism based on the language groups, and then distills the ultimate multilingual model from those assistants in an adaptive way. Experimental results derived from the TED dataset with 53 languages demonstrate the effectiveness of our approach in avoiding the negative transfer effect in MNMT, leading to an improved translation performance (about 1 BLEU score on average) compared to strong baselines.

1 Introduction

The surge over the past few decades in the number of languages used in electronic texts for international communications has promoted Machine Translation (MT) systems to shift towards multilingualism. However, most successful MT applications, i.e., Neural Machine Translation (NMT) systems, usually rely on supervised deep learning, which is notoriously data-hungry (Koehn and Knowles, 2017). Despite decades of research, high-quality annotated MT resources are only available

for a subset of the world’s thousands of languages (Paolillo and Das, 2006). Hence, data scarcity is one of the significant challenges which comes along with the language diversity and multilingualism in MT. One of the most widely-researched approaches to tackle this problem is unsupervised learning which takes advantage of available unlabeled data in multiple languages (Lample et al., 2017; Arivazhagan et al., 2019; Snyder et al., 2010; Xu et al., 2019). However, unsupervised approaches have relatively lower performance compared to their supervised counterparts (Dabre et al., 2020). Nevertheless, the performance of the supervised MNMT models is highly dependent on the types of languages used to train the model (Tan et al., 2019a). If languages are from very distant language families, they can lead to *negative transfer* (Torrey and Shavlik, 2010; Rosenstein, 2005), causing lower translation quality compared to the individual bilingual counterparts.

To address this problem, some improvements have been achieved recently with solutions that employ some sort of supervision to guide MNMT using *linguistic typology* (Oncevay et al., 2020; Chowdhury et al., 2020; Kudugunta et al., 2019; Bjerva et al., 2019). The linguistic typology provides this supervision by treating the world’s languages based on their functional and structural characteristics (O’Horan et al., 2016). Taking advantage of this property, which explains both language similarity and language diversity, we aim in our approach to combine two solutions for training an MNMT model: (a) creating a universal, language-independent MNMT model (Johnson et al., 2017); (b) systematically designing the possible variations of language-dependent MNMT models based on the language relations (Maimaiti et al., 2019).

Our approach to preventing negative transfer in MNMT is to group models which behave similarly in separate language clusters. Then, we perform a Knowledge Distillation (KD) (Hinton et al.,

*Corresponding author

2015) approach by selectively distilling the bilingual teacher models’ knowledge in the same language cluster to a multilingual teacher-assistant model. The intermediate teacher-assistant models are representative of their own language cluster. We further adaptively distill knowledge from the multilingual teacher-assistant models to the ultimate multilingual student. In summary, our main contributions are as follows:

- We use *cluster-based teachers* in a hierarchical knowledge distillation approach to prevent negative transfer in MNMT. Different from the previous cluster-based approaches in multilingual settings (Oncevay et al., 2020; Tan et al., 2019a), our approach makes use of all the clusters with a *universal* MNMT model while retaining the language relatedness structure in a hierarchy.
- We distill the ultimate MNMT model from multilingual teacher-assistant models, each of which represents one language family and usually perform better than the individual bilingual models from the same language family. Thus, the cluster-based teacher-assistant models can lead to a better knowledge distillation compared to a diverse set of bilingual teacher models as used in multilingual KD (Tan et al., 2019b).
- We explore *a mixture of linguistic features* by utilizing different clustering approaches to obtain the cluster-based teacher-assistants. As the language groups created by different language feature vectors can contribute differently to translation, we adaptively distill knowledge from teacher-assistant models to the ultimate student to improve the knowledge gap of the student.
- We perform extensive experiments on 53 languages, showing the effectiveness of our approach in avoiding negative transfer in MNMT, leading to an improved translation performance (about 1 BLEU score on average) compared to strong baselines. We also conduct comprehensive ablation studies and analysis, demonstrating the impact of language clustering in MNMT for different language families and in different resource-size scenarios.

2 Related Work

The majority of works on MNMT mainly focus on different architectural choices varying in the degree of parameter sharing in the multilingual setting. For example, the works based on the idea of minimal parameter sharing share either encoder, decoder,

or attention module (Firat et al., 2017; Lu et al., 2018), and those with complete parameter sharing tend to share entire models (Johnson et al., 2017; Ha et al., 2016). In general, these techniques implicitly assume that a set of languages is pre-given without considering the positive or negative effect of language transfer between the languages shared in one model. Hence, they can usually achieve comparable results with individual models (trained with individual language pairs) only when the languages are less diverse or the number of languages is small. When several diverse language pairs are involved in training an MNMT system, the *negative transfer* (Torrey and Shavlik, 2010; Rosenstein, 2005) usually happens between more distant languages, resulting in degraded translation accuracy in the multilingual setting. To address this problem, Tan et al. (2019a) suggested a clustering approach using either prior knowledge of language families or using language embedding. They obtained the language embedding by retrieving the representation of a language tag which is added to the input of an encoder in a universal MNMT model. Later, Oncevay et al. (2020) introduced another clustering technique using the multi-view language representation. They fused language embeddings learned in an MNMT model with syntactic features of a linguistic knowledge base (Dryer and Haspelmath, 2013). Tan et al. (2019b) proposed a knowledge distillation approach which transfers knowledge from bilingual teachers to a multilingual student when the accuracy of teachers are higher than the student. Their approach eliminates the accuracy gap between the bilingual and multilingual NMT models. However, we argue that distilling knowledge from a *diverse* set of parent models into a student model can be sub-optimal, as the parents may compete instead of collaborating with each other, resulting in negative transfer.

3 Hierarchical Knowledge Distillation

We address the problem of *data scarcity* and *negative transfer* in MNMT with a *Hierarchical Knowledge Distillation* (HKD) approach. The hierarchy in HKD is constructed in such a way that the node structure captures the similarity structure and the relatedness of the languages. Specifically, in an inverse pyramidal structure as shown in Figure 1, the root node corresponds to the ultimate MNMT model that we aim to train, the leave nodes correspond to each individual bilingual NMT models,

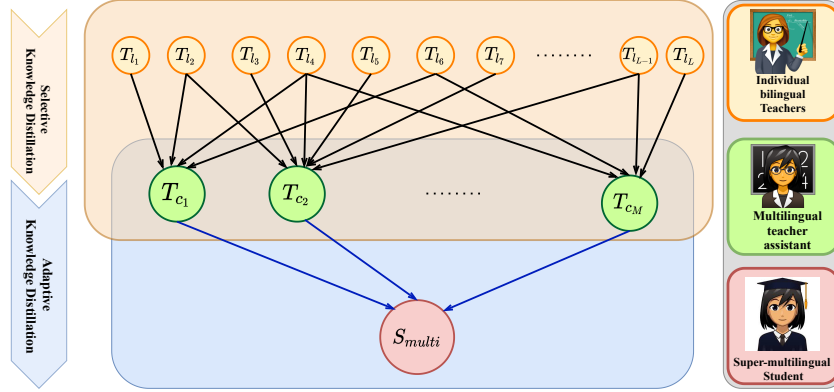


Figure 1: **HKD** approach: In the first phase of knowledge distillation, aka “Selective KD”, the knowledge is transferred from bilingual teacher models per clusters (orange circles) to the multilingual teacher-assistant models (green circles). For example $T_{l_1}, T_{l_2}, T_{l_4}$, and T_{l_6} are belonged to one cluster and distilled to teacher-assistant model T_{c_1} . In the second KD phase, aka “Adaptive KD”, knowledge is transferred from ensemble of intermediate related teacher-assistant models to the ultimate student (red circle) adaptively.

and the non-terminal nodes represent the language clusters. Our hypothesis is that leveraging common characteristics of languages in the same language group, which is formed using clustering algorithms based on the typological properties of languages (O’Horan et al., 2016), the HKD method can train a high quality MNMT model by distilling knowledge from related languages, rather than diverse ones.

Our HKD approach consists of two knowledge distillation mechanisms, providing two levels of supervision for training the ultimate MNMT model (illustrated in Figure 1), including: (i) **selective distillation** of knowledge from individual bilingual teachers to the multilingual intermediate teacher-assistants, each of which corresponds to one language group; and (ii) **adaptive distillation** of knowledge from all related cluster-wise teacher-assistants to the super-multilingual ultimate student model in each mini-batch of training per language pair. Note that we do not utilize multilingual adaptive KD in both distillation phases as we need to have the predictions of all the *relevant experts* in adaptive KD. Using adaptive KD for both stages is particularly impractical when there is a huge set of diverse teachers as in the first phase. Hence, in the first distillation phase, we aim to generate the cluster-wise teacher assistants using selective KD as the pre-requisites for the adaptive KD phase. The main steps of HKD are elaborated as follows:

Clustering: Clustering can be conducted using different language vectors such as: i) sparse language vectors from typological knowledge base (KB) databases, ii) dense learned language embedding vectors from multilingual NLP tasks, and iii) the combination of KB and task-learned language vec-

tors. The implicit causal relationships between languages are usually learned from translation tasks; the genetic, the geographical, and the structural similarities between languages are extracted from typological KBs (Bjerva et al., 2019). Thus, the language groups created by different language vectors can contribute differently to the translation and it is not quite clear which types of language features are more helpful in MNMT systems (Oncevay et al., 2020). For example, “Greek” can be clustered with “Arabic” and “Hebrew” based on the mix of KB and task-learned language vectors. Meanwhile, it can be clustered with “Macedonian” and “Bulgarian” based on NMT-learned language vectors. Therefore, we cluster the languages based on all types of language representations and propose to explore a mixture of linguistic features by utilizing all clusters in training the ultimate MNMT student. So, given a training dataset consisting of L languages and K clustering approaches, where each clustering approach creates n clusters, we are interested in training a many-to-one MNMT model (ultimate student) by hierarchically distilling knowledge from all M clusters to the ultimate student, where $M := \sum_{k=1}^K n_k$.

Multilingual selective knowledge distillation:

Assume we have a language cluster that consists of L' languages, where $l \in \{1, 2, \dots, L'\}$. Given a collection of pretrained individual teacher models $\{\theta^l\}_{l=1}^{L'}$, each handling one language pair in $\{\mathcal{D}^l\}_{l=1}^{L'}$, and inspired by Tan et al. (2019b), we use the following knowledge distillation objective

for each language l in the cluster.

$$\mathcal{L}_{KD}^{selective}(\mathcal{D}^l, \theta^c, \theta^l) := - \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{D}^l} \sum_{t=1}^{|\mathbf{y}|} Q(v|\mathbf{y}_{<t}, \mathbf{x}, \theta^l) \log P(v|\mathbf{y}_{<t}, \mathbf{x}, \theta^c) \quad (1)$$

where θ^c is the teacher assistant model, $|V|$ is the vocabulary set, $P(\cdot | \cdot)$ is the conditional probability of the teacher assistant model, and $Q(\cdot | \cdot)$ denotes the output distribution of the bilingual teacher model. According to Eq. (1), knowledge distillation regularises the predictive probabilities generated by a cluster-wise multilingual model with those generated by each individual bilingual models. Together with the translation loss (\mathcal{L}_{NLL}), we have the following selective KD loss to generate the intermediate teacher-assistant model:

$$\mathcal{L}_{ALL}^{selective}(\mathcal{D}^l, \theta^c, \theta^l) := (1 - \lambda)\mathcal{L}_{NLL}(\mathcal{D}^l, \theta^c) + \lambda\mathcal{L}_{KD}^{selective}(\mathcal{D}^l, \theta^c, \theta^l) \quad (2)$$

where λ is a tuning parameter that balances the contribution of the two losses. Instead of using all language pairs in Eq (1), we used a deterministic but dynamic approach to exclude language pairs from the loss function if the multilingual student surpasses the individual models on some language pairs during the training, which makes the training selective. This selective distillation process¹ is applied to all clusters obtained from different clustering approaches. It is noteworthy that (i) the selective knowledge distillation generates a teacher-assistant model for each cluster, i.e., $c \in \{1, 2, \dots, M\}$; (ii) each language can be in multiple clusters due to the use of different language representations, thus there can be more than one effective teacher-assistant model for any given language pair (illustrated in Figure 2.). So for each language pair, we have a set of effective clusters: $c \in \{1, 2, \dots, C_{sim}\}$.

Multilingual adaptive knowledge distillation: Given a collection of effective teacher-assistant models $\{\theta^c\}_{c=1}^{C_{sim}}$, where C_{sim} is the number of effective clusters per language, we devise the fol-

¹The training algorithm of selective knowledge distillation is summarized in Alg. 1 in (section A.1 - Appendix), which is similar to the one used in (Tan et al., 2019b).

Algorithm 1: Multilingual Adaptive KD

Input : Training corpora: $\{\mathcal{D}^l\}_{l=1}^L$, where
 $\mathcal{D}^l := \{(\mathbf{x}_1^l, \mathbf{y}_1^l), \dots, (\mathbf{x}_n^l, \mathbf{y}_n^l)\}$;
List of languages: L ;
List of language clusters: $\{C^m\}_{m=1}^M$;
Cluster-based MNMT models: $\{\theta^c\}_{c=1}^M$;
Total training epochs: N ;
Batch size: J ;

Output : Ultimate multilingual student model: θ_s ;
Randomly initialize multilingual model θ_s , accumulated gradient
 $g = 0$, distillation flag $f^l = True$ for $l \in L$;
 $n = 0$;
while $n < N$ **do**
 $g = 0$;
 $C_{sim} = []$;
for $l \in L$ **do**
// find the effective clusters with similar languages;
for $c \in \{C\}_1^M$ **do**
if $l \in c$ **then**
| $C_{sim}.append(c)$
 $D^l = random_permute(\mathcal{D}^l)$;
 $\mathbf{b}_1^l, \dots, \mathbf{b}_J^l = create_minibatches(\mathcal{D}^l)$
//where $\mathbf{b}^l = (\mathbf{x}^l, \mathbf{y}^l)$;
 $j = 1$;
while $j \leq J$ **do**
// compute contribution weights;
for $c \in C_{sim}$ **do**
| $\Delta_c = -ppl(\theta^c(\mathbf{b}_j^l))$;
 $\alpha = softmax(\Delta_1, \dots, \Delta_c)$;
//compute the gradient on loss $\mathcal{L}_{ALL}^{adapt.}$;
 $\mathbf{g} = \nabla_{\theta_s} \mathcal{L}_{ALL}^{adapt.}(\mathbf{b}_j^l, \theta_s, \{\theta^c\}_1^C, \alpha)$;
// updates the parameters;
 $\theta_s = update_param(\theta_s, \mathbf{g})$;
 $j = j + 1$;
 $n = n + 1$;

lowing KD objective for each language pair,

$$\mathcal{L}_{KD}^{adaptive}(\mathcal{D}^l, \theta_s, \{\theta^c\}_1^{C_{sim}}, \alpha) := - \sum_{c=1}^{C_{sim}} \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{D}^l} \alpha_c \sum_{t=1}^{|\mathbf{y}|} \sum_{v \in V} Q(v|\mathbf{y}_{<t}, \mathbf{x}, \theta^c) \log P(v|\mathbf{y}_{<t}, \mathbf{x}, \theta_s) \quad (3)$$

where α dynamically weigh the contribution of the teacher-assistants/clusters. α is computed via an attention mechanism based on the rewards (negative perplexity) attained by the teachers on the data, where these values are passed through a softmax transformation to turn into a distribution (Saleh et al., 2020). This adaptive distillation of knowledge allows the student model to get the best of teacher-assistants (which are representative of different linguistic features) based on their effectiveness to improve the knowledge gap of the student. The total loss function then becomes a weighted combination of losses coming from the ensemble of teachers and the data,

$$\mathcal{L}_{ALL}^{adaptive}(\mathcal{D}^l, \theta_s, \{\theta^c\}_1^{C_{sim}}, \alpha) := \lambda_1 \mathcal{L}_{NLL}(\mathcal{D}^l, \theta_s) + \lambda_2 \mathcal{L}_{KD}^{adapt.}(\mathcal{D}^l, \theta_s, \{\theta^c\}_1^{C_{sim}}, \alpha) \quad (4)$$

cluster 1	cluster 2	cluster 3	cluster 4	cluster 5	cluster 6	cluster 7	cluster 8	cluster 9	cluster 10	
Japanese Korean Mongolian Burmese	Malay Thai Chinese Indonesian Vietnamese	Marathi Tamil Bengali Georgian	Kurdish Basque Persian Hindi Kazakh Urdu	Greek Arabic Hebrew	Turkish Azerbaijani Finnish	Hungarian Armenian	Slovak Polish Russian Macedonian Lithuanian Belarusian Ukrainian	Croatian Bosnian Slovenian Serbian Czech Estonian Albanian	Galician Italian Portuguese Bulgarian Romanian Spanish	French Dutch Danish German Swedish Bokmal

Table 1: **Clustering type (1)** – SVCCA-53 (Onccey et al., 2020), Based on multi-view representation using both syntax features of WALS and language vectors learned by MNMT model trained with TED-53.

cluster 1	cluster 2	cluster 3	cluster 4	cluster 5	cluster 6	cluster 7	cluster 8	cluster 9	cluster 10							
Korean Bengali Marathi	Hindi Urdu	Basque Arabic Hebrew	Armenian Persian Kurdish	Hungarian Azerbaijani Mongolian	Turkish Japanese	Georgian Tamil	Kazakh Burmese	Bosnian Albanian Polish Slovak	Croatian Macedonian Belarusian Estonian	Russian Ukrainian Slovenian Lithuanian	Finnish Czech Serbian	Thai Malay Indonesian Vietnamese Chinese	Romanian Spanish Italian Galician	Bulgarian Swedish German	Dutch Portuguese French	Danish Greek Bokmal

Table 2: **Clustering type (2)** – SVCCA-23 (Onccey et al., 2020), Based on multi-view representation using both syntax features of WALS and language vectors learned by MNMT model trained with WIT-23.

cluster 1	cluster 2	cluster 3	cluster 4	cluster 5	cluster 6	cluster 7	cluster 8	cluster 9	cluster 10	cluster 11			
Estonian Finnish	Hindi Burmese Armenian Georgian	Basque Belarusian Bengali Kurdish Bosnian Marathi	Urdu Azerbaijani Mongolian Kazakh Malay	Galician French Italian Spanish Portuguese	Bokmal Danish Swedish	German Dutch	Chinese Japanese Korean Hungarian Turkish	Slovak Polish Russian Ukrainian	Lithuanian Slovenian Croatian Serbian Czech	Persian Indonesian Hebrew Vietnamese	Thai Arabic	Romanian Albanian	Macedonian Bulgarian Greek

Table 3: **Clustering type (3)** – Based on NMT-learned representation using a set of 53 factored language embeddings (Onccey et al., 2020; Tan et al., 2019a).

cluster 1	cluster 2	cluster 3
Bengali Chinese Korean	Kazakh Azerbaijani Mongolian Marathi	Burmese Urdu Turkish Tamil Hindi
Thai Vietnamese Indonesian Malay	Armenian Georgian Macedonian	Bosnian Slovenian Hungarian
	Basque Serbian Dutch	Albanian Greek German
	Galician Portuguese	Croatian Spanish
	French Lithuanian	Danish Hebrew
	Polish Russian	Finnish Kurdish
	Italian Estonian	Arabic Swedish

Table 4: **Clustering type (4)** – Based on KB representation using syntax features of WALS (Onccey et al., 2020).




Figure 2: Effective teachers for each language after clustering. C refers to the clustering type and T refers to the Teacher. For language a , we have two *effective* teachers: $T_{C_1}^1$ and $T_{C_2}^1$.

The training process is summarized in Alg. 1.

4 Experiment

In this section we explain our experiment settings as well as our experimental findings.

Data: We conducted extensive experiments on a parallel corpus (53 languages \rightarrow English) from TED talks transcripts² created and tokenized by Qi et al. (2018). This corpus has 26% of language pairs having less than or equal to 10k sentences (extremely low-resource), and 33% of language pairs having less than 20k sentences (low-resource). All the sentences were segmented with BPE segmentation (Sennrich et al., 2016) with a learned BPE model with 32k merge operations on all languages. We

²<https://github.com/neulab/word-embeddings-for-nmt>

kept the output vocabulary of the teacher and student models the same to make the knowledge distillation feasible. Details about the size of training data, language codes, and preprocessing steps are described in Section A.2.

Clustering: We clustered all the languages based on the three different types of representations discussed in Section 3 in order to take advantage of a mixture of linguistic features while training the ultimate student. Following Onccey et al. (2020), we adopted their multi-view language representation approach that uses Singular Vector Canonical Correlation Analysis – SVCCA (Raghu et al., 2017) to fuse the one-hot encoded KB representation obtained from syntactic features of WALS (Dryer and Haspelmath, 2013) and a dense NMT-learned view obtained from MNMT (Tan et al., 2019a). Specifically, SVCCA-53 uses 53 languages of TED dataset to build the language representations and generates 10 clusters, the languages within each of which usually have the same *phylogenetic* or *geographical* features. SVCCA-23 instead uses 23 languages of WIT-23 (Cettolo et al., 2012) to compute the shared space. We also generated language clusters based on either KB-based representation using syntax features of WALS (Dryer and Haspelmath, 2013) and NMT-learned representation alone. Tables 1-4 show the generated language clusters.

IE/Balto-Slavic	IE/Italic	IE/Indo-Iranian	IE/Germanic	Turkic	Uralic	Afroasiatic	Sino-Tibetan	Austroasiatic	Koreanic	Japonic	Austrasiatic	IE/Hellenic	Kra-Dai	IE/Albanian	IE/Armenian	Kartvelian	Mongolic	Dravidian	Isolate (Basque)
be, bs, sl, mk, lt, sk cs, uk, hr, sr, bg, pl, ru	bn, ur, ku hi, fa, nr	gl, pt, ro fr, es, it	nb, da sv, de, nl	kk, az, tr	et, fi, hu	he, ar	my, zh	ms, id	ko	ja	vi	el	th	sq	hy	ka	mn	ta	eu

Table 5: Language families (Eberhard et al., 2019). IE refers to Indo European.

Training Configuration: All models were trained with Transformer architecture (Vaswani et al., 2017). The individual baseline models were trained with the model hidden size of 256, feed-forward hidden size of 1024, and 2 layers. All multilingual models were trained with the model hidden size of 512, feed-forward hidden size of 1024, and 6 layers. We used selective multilingual knowledge distillation approach for training our cluster-based MNMT models as selective multilingual KD perform better than universal MNMT without KD (Tan et al., 2019b). We did not carry out intense parameter tuning to search for the best parameter settings of each model for the sake of simplicity. It is noteworthy that the purpose of our experiments is rather to demonstrate the benefit of considering different type of language clusters through a unified hierarchical knowledge distillation. Details about the training configuration is provided in Section A.2.

4.1 Experimental Results

The translation results of (53 languages → English) for all approaches are summarised in Table 6. The language pairs are sorted based on the size of training data in an ascending order. The translation quality is evaluated and reported based on the BLEU score (Papineni et al., 2002).

4.1.1 Studies of cluster-based MNMT models

In this section, we discuss the cluster-based MNMTs’ results through the following observations.

Low resource vs high resource languages: All cluster-based MNMT and baseline approaches are ranked based on the number of the times they got the first or second-best score in different resource-size scenarios in Table 7. Based on this result and also the result represented in Table 6, the massive MNMT models with all languages (second column under baseline and first column under selective KD in Tables 6, 7) outperform the cluster-based MNMT models (columns (2-5) under selective KD in Tables 6, 7) in extremely low-resource scenarios (e.g., bn-en, ta-en, eu-en). This result shows that having more data either from related languages or distant

Resource	src	size #sent. (k)	Baseline		Multilingual Selective KD				Multi. HKD	
			Individ.	Universal Multi.	All langs	Clus. type1	Clus. type2	Clus. type3		Clus. type4
Extremely low resource	kk	3.3	3.42	5.05	4.66	<u>7.00</u>	3.51	3.10	8.13	6.61
	be	4.5	5.13	12.51	12.36	12.78	10.81	8.46	15.18	<u>14.88</u>
	bn	4.6	5.06	12.50	<u>12.58</u>	9.13	10.11	12.16	10.13	12.81
	eu	5.1	4.4	13.12	12.00	9.08	9.70	8.14	11.03	<u>12.90</u>
	ms	5.2	3.78	13.88	14.61	12.93	12.94	7.63	12.98	<u>14.11</u>
	bs	5.6	7.92	14.82	15.46	<u>18.05</u>	16.89	9.03	19.02	17.51
	az	5.9	5.79	<u>10.32</u>	9.91	9.59	9.17	8.64	9.23	10.40
	ur	5.9	8.98	12.76	16.50	13.35	13.17	12.02	13.39	<u>15.95</u>
	ta	6.2	4.57	<u>5.86</u>	6.19	4.02	5.76	3.96	3.49	5.63
	mn	7.6	3.54	6.11	5.82	5.75	<u>6.60</u>	5.39	6.20	6.81
	mr	9.8	6.92	10.53	<u>10.72</u>	8.70	9.00	8.39	9.04	10.98
	gl	10.0	13.5	22.04	22.44	25.53	26.81	<u>26.93</u>	25.90	27.11
	ku	10.3	6.3	10.32	12.12	9.43	6.98	9.22	<u>12.93</u>	13.03
	et	10.7	8.24	12.21	13.19	13.47	13.21	10.44	<u>13.94</u>	14.10
Low resource	ka	13.1	8.64	8.18	8.66	9.28	<u>10.85</u>	9.14	8.88	11.15
	nb	15.8	26.36	28.49	29.08	33.55	34.31	28.79	30.87	<u>33.89</u>
	hi	18.7	10.66	16.03	17.93	13.27	12.09	12.16	12.80	<u>16.11</u>
	sl	19.8	11.45	15.12	15.39	16.48	16.54	17.75	<u>18.31</u>	18.43
	Enough resource	hy	21.3	11.14	14.07	15.12	13.76	12.77	10.81	17.17
my		21.4	4.91	<u>10.70</u>	11.11	9.65	6.35	8.48	9.54	8.81
fi		24.2	8.16	11.69	12.23	11.36	12.57	10.59	<u>12.76</u>	12.90
mk		25.3	18.32	20.63	21.09	21.48	20.06	<u>24.65</u>	23.8	25.05
lt		41.9	14.78	15.44	16.76	16.98	16.9	17.96	18.24	<u>18.11</u>
sq		44.4	22.62	24.44	25.22	24.74	23.22	<u>26.89</u>	26.42	26.93
da		44.9	31.85	30.39	30.61	35.02	39.76	30.58	32.04	<u>36.00</u>
sv		56.6	27.2	27.18	26.84	31.36	34.52	26.81	28.65	<u>33.14</u>
sk		61.4	19.36	22.04	22.58	22.49	21.18	24.08	23.77	24.33
id		87.4	20.51	20.89	20.69	21.11	21.13	<u>22.56</u>	21.12	22.76
tr		96.9	20.46	21.34	21.72	22.94	22.94	<u>23.09</u>	22.87	23.30
cs		103.0	20.13	22.01	22.07	21.72	22.49	<u>23.12</u>	22.86	23.62
uk		108.4	21.32	22.11	23.07	23.06	22.91	23.58	<u>23.66</u>	24.09
hr		122.0	25.89	26.51	27.17	27.56	25.62	<u>28.66</u>	28.34	28.91
el		134.3	26.82	26.07	28.51	30.05	31.35	29.13	29.66	<u>30.10</u>
sr		136.8	26.94	25.43	25.88	27.48	25.75	<u>27.69</u>	27.12	27.97
hu		147.1	18.46	17.61	18.55	19.08	18.41	20.16	19.82	<u>20.10</u>
fa		150.8	23.60	21.7	21.29	21.31	22.44	<u>23.51</u>	22.24	23.19
de		167.8	15.23	14.83	16.69	16.88	17.79	15.44	16.67	18.04
ja		168.2	10.11	8.61	8.93	10.14	10.14	<u>10.17</u>	8.69	10.30
vi		171.9	18.97	19.19	20.58	21.60	21.60	<u>21.30</u>	20.33	21.82
bg		174.4	28.85	27.66	29.14	31.67	<u>32.18</u>	29.86	30.48	32.33
pl		176.1	17.23	18.62	19.45	19.45	18.37	19.93	<u>20.26</u>	20.71
ro		180.4	25.21	25.97	26.53	28.03	28.43	<u>28.90</u>	27.35	29.00
tr		182.3	17.72	10.2	10.01	<u>18.66</u>	18.19	19.85	18.27	16.91
nl		183.7	27.65	26.91	26.82	28.05	29.07	29.17	28.03	29.58
zh		184.8	20.44	22.10	22.71	22.19	22.11	23.91	22.84	<u>23.85</u>
es	195.9	30.17	29.55	30.00	29.06	33.45	31.46	31.82	<u>32.76</u>	
it	204.4	26.84	25.13	27.99	30.57	<u>30.85</u>	30.36	29.45	30.93	
ko	205.4	15.98	15.71	16.41	15.17	15.18	<u>17.45</u>	16.00	17.70	
ru	208.4	19.76	19.83	20.86	20.85	20.80	21.25	<u>21.49</u>	21.77	
he	211.7	29.35	28.03	28.27	<u>32.82</u>	<u>32.18</u>	31.32	30.02	32.05	
fr	212.0	30.08	30.55	30.28	<u>32.25</u>	32.19	31.14	31.34	32.65	
ar	213.8	25.36	23.89	24.48	<u>28.53</u>	28.03	27.46	25.85	28.94	
pt	236.4	30.99	31.12	30.85	33.36	<u>33.84</u>	33.25	32.56	33.90	
Avg.	-	-	18.50	18.64	19.24	19.84	19.87	19.35	<u>20.05</u>	21.16

Table 6: BLEU scores of the translation tasks for 53 Languages→English. The bold numbers show the best scores and the underline ones show the 2nd best results. “All langs” refers to all 53 languages trained with a multilingual selective KD.

languages has the most impact on training a better MNMT model for under-resourced languages. Furthermore, clustering type (4) is dominant among other clustering approaches for under-resourced situations. This result is also explainable based on

the size of the clusters in clustering type (4). The translations of extremely low resource languages are significantly improved when they have been clustered in the third cluster of clustering type (4) with 35 languages (shown in Table 4). However, for languages with enough resources, the multilingual baselines with all languages under-performed other cluster-based MNMT models.

Related vs isolated languages: A group of languages that originated from a similar ancestor is known as a *language family*; and a language that does not have any relationship with another languages is called a *language isolate*. The language families (Eberhard et al., 2019) are shown in Table 5. According the results shown in Table 6, clustering approaches usually have the same behaviour and less diversity for clustering languages belonged to IE/Germanic, IE/Italic, Afroasiatic, and Austronesian families. In comparison, there is more diversity, and less consensus for clustering languages belonged to IE/Balto-Slavic, IE/Indo-Iranian, Turkic, Uralic, and Sino-Tibetan families. Moreover, the isolated languages (shown in the last 11 columns of Table 5) generally have the same behaviour and less variance in BLEU scores in different clustering approaches. This observation shows that cluster-based MNMT models (regardless of the clustering type) do not significantly improve the translation of isolated languages unless the isolated languages have extremely low resources and have been clustered in a huge cluster (e.g., eu-en, hy-en). The results of cluster-based MNMT are presented based on the language families in Section A.3.

Random clustering vs Actual clustering: We conducted an ablation study by using clusters with randomly chosen languages for two translation tasks (el→en and gl→en). We kept the number of languages per cluster the same as the actual clustering to make a fair comparison. According to the result represented in Table 9, for both translation tasks, random clusters underperform the actual clusters in all clustering types. However, notably, the average BLEU score’s difference between the random and actual clusters for gl→en is considerably higher than el→en (gl→en, $\Delta = -8.14$ vs el→en, $\Delta = -1.9$). This observation is inline with the previous observation that *Greek* (el) is an isolated language categorised in IE/Hellenic family and clustering approaches have less impact on this language due to its lower similarity to most languages. In comparison, *Galician* (gl) is highly sim-

ilar to the languages in IE/Italic family and clustering improves the translation of gl→en remarkably. We compared the translation results between actual cluster-based fake cluster-based multilingual approaches in Section A.3.

4.1.2 Studies of HKD

According to the results in Table 6, the HKD approach outperforms massive and cluster-based MNMT models in average by 1.11 BLEU score. We discuss it more in the following observations:

Ranking based on data size: We ranked all approaches, including HKD, based on the number of times they got the first or second-best score (shown in Table 8). According to this result, the HKD approach in three different situations, i.e. extremely low resource, low resource, and enough resource, has the best rank among other approaches. This observation proves that the HKD approach is robust in different *data-size* situations by leveraging the best of both multilingual NMT and language-relatedness guidance in a systematic HKD setting.

Clustering consistency impact on HKD: Based on the result in Table 6, the HKD approach underperforms other multilingual approaches when the clusters are inconsistent, causing a high variance in teacher-assistants’ results. For example for kk→en, bs→en, the variance of the BLEU scores of the teacher-assistant models is 6.92 and 20.81 respectively and HKD underachieved a good result. This observation shows that, although in the second phase of the HKD approach, the cluster-based teacher-assistants adaptively contribute to training the ultimate students, still a weak teacher-assistant deteriorates the collaborative teaching process. One possible solution is excluding the worst teacher-assistant in such heterogeneous situations.

4.1.3 Comparison with other approaches

To highlight the pros and cons of the related baselines (with and without KD), we draw a comparison shown in Table 10. Accordingly, our HKD approach is comparable with other approaches based on the following properties:

Multilingual translation: Our approach works in a multilingual setting by sharing resources between high-resource and low-resource languages. This property not only improves the regularisation of the model by avoiding over-fitting to the limited data of the low-resource languages but also decreases the deployment footprint by leveraging the whole training in a single model instead of having individual

Resource-size	size (# sent.)	Baseline		Multilingual Selective KD				
		Individ.	Multi.	All langs	Clus. type1	Clus. type2	Clus. type3	Clus. type4
Extremely low resource	<= 10k	0%	21.43%	28.57%	14.28%	7.14%	3.58%	<u>25.00%</u>
Low resource	> 10k and <= 20k	0%	12.50%	12.50%	25.00%	25.00%	12.50%	12.50%
Enough resource	> 20k	1.39%	1.39%	2.78%	20.83%	<u>25.00%</u>	27.78%	20.83%

Table 7: The translation ranking ablation study for all approaches excluding the HKD approach based on the percentage of the times they got the 1st or 2nd best results. Sum of percentages in each row = 100%.

Resource-size	size (# sent.)	Baseline		Multilingual Selective KD					HKD
		Individ.	Multi.	All langs	Clus. type1	Clus. type2	Clus. type3	Clus. type4	
Extremely low resource	<= 10k	0%	13.79%	17.24%	6.90%	3.45%	3.45%	<u>17.24%</u>	37.93%
Low resource	> 10k and < 20k	0%	0%	12.50%	0%	<u>25.00%</u>	0%	12.50%	50.00%
Enough resource	>20k	1.43%	1.43%	1.43%	5.71%	12.86%	<u>24.28%</u>	8.57%	44.29%

Table 8: The ranking of all approaches based on the percentage of the times they got the 1st or 2nd best results.

Model	Contrib. Langs	BLEU	Contrib. Langs	BLEU
Clus. Rand.1	gl, nb, uk, hr, se, ja	16.53	el, id, be	28.01
Clus. Rand.2	gl, ta, mk, be, id, sq, pt, fr, ur, az, ku, bs, fa	20.27	el, cs, lt, id, sk, th, it, hy, ms, hu, mk, my, bn	27.78
Clus. Rand.3	gl, az, ja, nb, kk	13.61	el, sq, th	27.97
Clus. Rand.4	gl, zh, pt, fa, ar, kk, sr, bg, nl, cs, th, ko, vi, hu, mk, fi, ru, mn, de, sl, el, ka, pl, et, ta, fr, ur, ro, sv, mr, be, bs, uk, sq, az	22.20	el, zh, pt, fa, ar, kk, sr, bg, nl, cs, th, ko, vi, hu, mk, fi, ru, mn, de, sl, gl, ka, pl, et, ta, fr, ur, ro, sv, mr, be, bs, uk, sq, az	28.82
Avg.	-	18.15	-	28.14

Table 9: Ablation study on using random languages in all clustering types for (gl→en) and (el→en). The results for actual clustering are provided in Table 6. The average result of actual clustering for (gl→en) and (el→en) are 26.29 and 30.04 respectively.

models per language (Dabre et al., 2020).

Optimal transfer: In the HKD approach, we have an optimal transfer by transferring knowledge from *all* possible languages *related* to a student in the hierarchical structure which leads to the best average BLEU score (21.16) comparing to the other baselines (shown in Table 6). In the universal multilingual NMT without KD, the language transfer stream is maximized when all languages shared their knowledge in a single model during training; however, it is not an optimal transfer due to the lack of any condition on the relatedness of languages contributing in the multilingual training. The related experiments are shown in the second column under the baseline experiments in Table 6. The average BLEU score of this approach is 18.46. In multilingual selective KD (Tan et al., 2019b), knowledge is distilled from one selected teacher with the same language when training the student multilingually. So, although there is a condition on language relatedness, knowledge transfer

is not maximized as the similar languages from the same language family are ignored in the distillation process. The related results are shown in the first column under multilingual selective KD of Table 6. Accordingly, this approach got the average BLEU score of 19.24 in our experiments. Adaptive KD (Saleh et al., 2020) is a bilingual approach and also uses a random set of teachers which does not essentially have all the related languages to the student and does not lead to optimal transfer. We did not perform any experiment on adaptive KD (Saleh et al., 2020) since this is a bilingual approach.

Adaptive KD vs Selective KD vs HKD: All KD-based approaches in our comparison reduce *negative transfer* in different ways. In multilingual selective KD (Tan et al., 2019b), the risk of negative transfer is reduced by distilling knowledge from the selected teacher per language in a multilingual setting. In bilingual adaptive KD (Saleh et al., 2020), the contribution weights of different teachers vary based on their effectiveness to improve the student which prevents the negative transfer in bilingual setting. In HKD, the hierarchical grouping based on the language similarity provides a systematic guide to prevent negative transfer as much as possible. This property leads HKD to get the *best* results for 32 language pairs out of total 53 language pairs in our multilingual experiments.

5 Conclusion

We presented a Hierarchical Knowledge Distillation (HKD) approach to mitigate the negative transfer effect in MNMT when having a diverse set of languages in training. We put together all languages which behave similarly in the first phase of distillation process and generated the expert teacher-assistants for each group of languages. As

	Individual NMT	Uniform MNMT	Selective KD MNMT	Adaptive KD NMT	HKD MNMT
Multilingual	✗	✓	✓	✗	✓
Maximum transfer	●	✓	✗	✗	✓
KD from multiple languages	●	●	✗	✓	✓
Reduced risk of negative transfer	●	✗	✓	✓	✓

Table 10: Comparing different properties of HKD with: transformer-based individual and multilingual NMT (Vaswani et al., 2017), multilingual selective KD (Tan et al., 2019b), and adaptive KD (Saleh et al., 2020).

we clustered languages based on four different language representations capturing different linguistic features, we then adaptively distill knowledge from all related teacher-assistant models to the ultimate student in each mini-batch of training per language. Experimental results on 53 languages to English show our approach’s effectiveness to reduce negative transfer in MNMT. Our proposed approach is generalizable to one-to-many with the same setting as a many-to-one task. The clustering needs to be done in the target language, though. For many-to-many tasks, the hierarchical KD can be effective if the clustering is applied to both source and target languages. Our intended use, however, is many to one or one to many.

As the future direction, it is interesting to study an end-to-end HKD approach by adding a backward HKD pass compared to the forward HKD pass described in this paper.

Acknowledgements

This work was supported by the Multimodal Australian ScienceS Imaging and Visualisation Environment (MASSIVE) (www.massive.org.au).

References

Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2019. [The missing ingredient in zero-shot neural machine translation](#). *arXiv preprint arXiv:1903.07091*.

Johannes Bjerva, Robert Östling, Maria Han Veiga, Jörg Tiedemann, and Isabelle Augenstein. 2019. [What do language representations really represent?](#) *Computational Linguistics*, 45(2):381–389.

Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. [Wit3: Web inventory of transcribed and translated talks](#). In *Conference of european association for machine translation*, pages 261–268.

Koel Dutta Chowdhury, Cristina España-Bonet, and Josef van Genabith. 2020. [Understanding translationese in multi-view embedding spaces](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6056–6062.

Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. [A survey of multilingual neural machine translation](#). *ACM Computing Surveys (CSUR)*, 53(5):1–38.

Matthew S. Dryer and Martin Haspelmath. 2013. [The world atlas of language structures](#).

DM Eberhard, GF Simons, and CD Fennig. 2019. [What are the largest language families](#). *Ethnologue: Languages of the World*.

Orhan Firat, Kyunghyun Cho, Baskaran Sankaran, Fatos T Yarman Vural, and Yoshua Bengio. 2017. [Multi-way, multilingual neural machine translation](#). *Computer Speech & Language*, 45:236–252.

Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2016. [Toward multilingual neural machine translation with universal encoder and decoder](#). *Institute for Anthropomatics and Robotics*, 2(10.12):16.

Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). In *NIPS Deep Learning and Representation Learning Workshop*.

Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.

Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). *ACL 2017*, page 28.

Sneha Kudugunta, Ankur Bapna, Isaac Caswell, and Orhan Firat. 2019. [Investigating multilingual nmt representations at scale](#). 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)*, pages 1565–1575.

Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’ Aurelio Ranzato. 2017. [Unsupervised machine translation using monolingual corpora only](#). *arXiv preprint arXiv:1711.00043*.

Yichao Lu, Phillip Keung, Faisal Ladhak, Vikas Bhardwaj, Shaonan Zhang, and Jason Sun. 2018. [A neural interlingua for multilingual machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 84–92.

- Mieradilijiang Maimaiti, Yang Liu, Huanbo Luan, and Maosong Sun. 2019. Multi-round transfer learning for low-resource nmt using multiple high-resource languages. *ACM Transactions on Asian and Low-Resource Language Information Processing (TAL-LIP)*, 18(4):1–26.
- Helen O’Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, and Anna Korhonen. 2016. Survey on the use of typological information in natural language processing. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1297–1308.
- Arturo Oncevay, Barry Haddow, and Alexandra Birch. 2020. Bridging linguistic typology and multilingual machine translation with multi-view language representations. *arXiv preprint arXiv:2004.14923*.
- John C Paolillo and Anupam Das. 2006. Evaluating language statistics: The ethnologue and beyond. *Contract report for UNESCO Institute for Statistics*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wj Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. When and why are pre-trained word embeddings useful for neural machine translation? In *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, New Orleans, Louisiana. ACL.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. SVCCA: singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6078–6087.
- Michael T Rosenstein. 2005. To transfer or not to transfer. In *NIPS-2005 Workshop on Transfer Learning*, volume 898, pages 1–4.
- Fahimeh Saleh, Wray Buntine, and Gholamreza Haffari. 2020. Collective wisdom: Improving low-resource neural machine translation using adaptive knowledge distillation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3413–3421.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. ACL.
- Benjamin Snyder et al. 2010. *Unsupervised multilingual learning*. Ph.D. thesis, Massachusetts Institute of Technology.
- Xu Tan, Jiale Chen, Di He, Yingce Xia, QIN Tao, and Tie-Yan Liu. 2019a. Multilingual neural machine translation with language clustering. 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)*, pages 962–972.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019b. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- Lisa Torrey and Jude Shavlik. 2010. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI global.
- 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 30*, pages 5998–6008. Curran Associates, Inc.
- Chang Xu, Tao Qin, Gang Wang, and Tie-Yan Liu. 2019. Polygon-net: A general framework for jointly boosting multiple unsupervised neural machine translation models. In *IJCAI*, pages 5320–5326.

Multilingual Neural Machine Translation: Can Linguistic Hierarchies Help? (Supplementary Material)

Fahimeh Saleh Wray Buntine Gholamreza Haffari Lan Du
first.last@monash.edu
Monash University

1 Appendix

1.1 Selective Knowledge Distillation

This first stage of knowledge distillation is the same as the algorithm which is proposed by (Tan et al., 2019) and is summarized in Algorithm 1. This process is applied to all clusters obtained from different clustering approaches, $\{C^m := \{l_1, \dots, l_{|L'|}\}\}_{m=1}^M$ where M refers to the number of clusters and L' refers to the number of languages in each cluster.

1.2 Experiment Settings

Data. We conducted the experiments on a parallel corpus from TED talks transcripts¹ on 53 languages to English created and tokenized by (Qi et al., 2018). Detail about the size of training data and language codes based on ISO 639-1 standard² are listed in Table 1 and visualised in Figure 1. We concatenated all data which have the Portuguese-related languages in the source (pt→en, pt-br→en). We also concatenated all data with French-related languages in the source (fr→en, fr-ca→en). We removed any sentences in the training data which has overlap with any of the test sets. For multilingual training, as a standard practice (Wang et al., 2020), we up-sampled the data of low-resource language pairs to make all language pairs having roughly the same size and adjust the distribution of training data.

Training configuration. All models are trained with Transformer architecture (Vaswani et al., 2017), implemented in the Fairseq framework (Ott et al., 2019). The individual models are trained with the model hidden size of 256, feed-forward hidden size of 1024, and 2 layers. All multilingual models either cluster-based or universal MNMT models with or without knowledge distillation were

¹<https://github.com/neulab/word-embeddings-for-nmt>

²http://www.loc.gov/standards/iso639-2/php/English_list.php

Algorithm 1: Multilingual Selective Knowledge Distillation (Tan et al., 2019)

```

Input      : Training corpora:  $\{\mathcal{D}^l\}_{l=1}^L$ ; where
               $\mathcal{D}^l := \{(x_1^l, y_1), \dots, (x_n^l, y_n)\}$ ;
              List of all languages:  $L$ ;
              Individual models  $\{\theta^l\}_{l=1}^{L'}$ ;
              List of language pairs per cluster:  $L'$ ;
              Total training epochs:  $N$ ;
              Distillation check step:  $\mathcal{N}_{check}$ ;
              Threshold of distillation accuracy:  $\mathcal{T}$ 

Output    :  $\theta^c$ : multilingual model for each cluster,

Randomly initialize multilingual model  $\theta^c$ , accumulated gradient
 $g = 0$ , distillation flag  $f^l = True$  for  $l \in L'$ ;
 $n = 0$ ;
while  $n < N$  do
   $g = 0$ ;
  for  $l \in L'$  do
     $D^l = random\_permute(\mathcal{D}^l)$ ;
     $b_1^l, \dots, b_J^l = create\_minibatches(D^l)$ , where
     $b^l = (x^l, y)$ ;
     $j = 1$ ;
    while  $j \leq J$  do
      if  $f^l == True$  then
        //compute and accumulate the gradient on
        loss  $\mathcal{L}_{ALL}^{selective}$ ;
         $g = \nabla_{\theta^c} \mathcal{L}_{ALL}^{selective}(b_j^l, \theta^c, \theta^l)$ ;
        // updates the parameters using the
        optimiser ADAM;
         $\theta^c = update\_param(\theta^c, g)$ ;
      else
        //compute and accumulate the gradient on
        loss  $\mathcal{L}_{NLL}$ ;
         $g = \nabla_{\theta^c} \mathcal{L}_{NLL}(b_j^l, \theta^c, \theta^l)$ ;
        // updates the parameters using the
        optimiser ADAM;
         $\theta^c = update\_param(\theta^c, g)$ ;
       $j = j + 1$ ;
    if  $N \% \mathcal{N}_{check} == 0$  then
      for  $l \in L'$  do
        if  $Accuracy(\theta^c) < Accuracy(\theta^l) + \mathcal{T}$  then
           $f^l = True$ 
        else
           $f^l = False$ 
       $n = n + 1$ ;

```

trained with the model hidden size of 512, feed-forward hidden size of 1024, and 6 layers. We use the Adam optimizer (Kingma and Ba, 2015) and an inverse square root schedule with warmup (maximum LR 0.0005). We apply dropout and label smoothing with a rate of 0.3 and 0.1 for bilingual

TED-53 Languages

Language name	Kazakh	Belarusian	Bengali	Basque	Malay	Bosnian
Code	kk	be	bn	eu	ms	bs
train-size (#sent(k))	3.3	4.5	4.6	5.1	5.2	5.6
Language name	Azerbaijani	Urdu	Tamli	Mongolian	Marathi	Galician
Code	az	ur	ta	mn	mr	gl
train-size (#sent(k))	5.9	5.9	6.2	7.6	9.8	10
Language name	Kurdish	Estonian	Georgian	Bokmal	Hindi	Slovenian
Code	ku	et	ka	nb	hi	sl
train-size (#sent(k))	10.3	10.7	13.1	15.8	18.7	19.8
Language name	Kurdish	Estonian	Georgian	Bokmal	Hindi	Slovenian
Code	ku	et	ka	nb	hi	sl
train-size (#sent(k))	10.3	10.7	13.1	15.8	18.7	19.8
Language name	Armenian	Burmese	Finnish	Macedonian	Lithuanian	Albanian
Code	hy	my	fi	mk	lt	sq
train-size (#sent(k))	21.3	21.4	24.2	25.3	41.9	44.4
Language name	Danish	Swedish	Slovak	Indonesian	Thai	Czech
Code	da	sv	sk	id	th	cs
train-size (#sent(k))	44.9	56.6	61.4	87.4	96.9	103
Language name	Ukrainian	Croatian	Greek	Serbian	Hungarian	Persian
Code	uk	hr	el	sr	hu	fa
train-size (#sent(k))	108.4	122	134.3	136.8	147.1	150.8
Language name	German	Japanese	Vietnamese	Bulgarian	Polish	Romanian
Code	de	ja	vi	bg	pl	ro
train-size (#sent(k))	167.8	168.2	171.9	174.4	176.1	180.4
Language name	Turkish	Dutch	Chinese	Spanish	Italian	Korean
Code	tr	nl	zh	es	it	ko
train-size (#sent(k))	182.3	183.7	184.8	195.9	204.4	205.4
Language name	Russian	Hebrew	French	Arabic	Portuguese	
Code	ru	he	fr	ar	pt	
train-size (#sent(k))	208.4	211.7	212	213.8	236.4	

Table 1: Bilingual resources of 53 Languages \rightarrow English from TED dataset. Language names, language codes based on ISO 639-1 standard³, and training size based on the number of sentences in bilingual resources are shown in this table.

and multilingual models respectively.

For the first phase of distillation, i.e., the multilingual selective KD, the distillation coefficient λ is equal to 0.6. In the second phase of distillation, i.e., the multilingual adaptive KD, we applied $\lambda_1 = 0.5$ and λ_2 is started from 0.5 and increased to 3 using the annealing function of (Bowman et al., 2016; Saleh et al., 2020). We train our final multilingual student with mixed-precision floats on up to 8 V100 GPUs for maximum 100 epochs (≈ 3

days), with at most 8192 tokens per batch and early stopping at 20 validation steps based on the BLEU score. The translation quality is also evaluated and reported based on the BLEU (Papineni et al., 2002) score⁴.

⁴SacreBLEU signature:BLEU+case.mixed+numrefs.1+smooth.exp+tok.none+version.1.3.1




Figure 1: Training size (based on the number of sentences) for TED-53 bilingual resources (Language→English)

Model	Contributed Langs	BLEU	Contributed Langs	BLEU
Individual	gl	13.50	el	26.82
Multi. (uniform)	All langs.	22.04	All langs.	26.07
Multi. (SKD)	All langs.	22.44	All langs.	28.51
Clus. type1	gl, bg, ro, es, it, pt	25.53	el, ar, he	30.05
Clus. type2	gl, bg, sv, da, nb, de, nl, ro, el, es, it, fr, pt	26.81	el, bg, sv, da, nb, de, nl, es, it, gl, fr, pt, ro	31.35
Clus. type3	gl, fr, it, es, pt	26.93	el, mk, bg	29.13
Clus. type4	gl, fa, ku, eu, hu, et, fi, hy, ka, ar, he, fr, cs, lt, de, nl, it, sv, da, nb, ru, be, pl, bg, sl, pt, sr, uk, hr, mk, sk, ro, sq, el, es	25.90	el, fa, ku, eu, hu, et, fi, hy, ka, ar, he, fr, cs, lt, de, nl, it, sv, da, nb, ru, be, pl, bg, sl, pt, sr, uk, hr, mk, sk, ro, sq, es, gl	29.66
Avg.	-	26.29	-	30.04
Clus. Rand.1	gl, nb, uk, hr, se, ja	16.53	el, id, be	28.01
Clus. Rand.2	gl, ta, mk, be, id, sq, pt, fr, ur, az, ku, bs, fa	20.27	el, cs, lt, id, sk, th, it, hy, ms, hu, mk, my, bn	27.78
Clus. Rand.3	gl, az, ja, nb, kk	13.61	el, sq, th	27.97
Clus. Rand.4	gl, zh, pt, fa, ar, kk, sr, bg, nl, cs, th, ko, vi, hu, mk, fi, ru, mn, de, sl, el, ka, pl, et, ta, fr, ur, ro, sv, mr, be, bs, uk, sq, az	22.20	el, zh, pt, fa, ar, kk, sr, bg, nl, cs, th, ko, vi, hu, mk, fi, ru, mn, de, sl, gl, ka, pl, et, ta, fr, ur, ro, sv, mr, be, bs, uk, sq, az	28.82
Avg.	-	18.15 Δ -8.14	-	28.14 Δ -1.9

Table 2: Ablation study on using random clusters. Comparison of the (gl→en) and (el→en) translation tasks between individual, massive multilingual, and clustering-based multilingual (for actual and random clusters) baselines.

1.3 Analysis

Random clustering vs Actual clustering. We studied the behaviour of random clusters compared to the actual clusters in Section 4.1.1. For more clarification about the mentioned discussion, you can refer to Table 2.

Language Family. A group of languages that originated from a similar ancestor is known as a *language family*; and a language that does not have

any relationship with another languages is called a *language isolate*. We already discussed the ablation study’s results regarding the language families in comparison with isolated languages in Section 4.1.1 of the paper. For better visualisation of results based on language family, we sorted all language pairs based on their language family relations in Tables 3, 4, 5, and 6.

Family	lang	Clustering Type 1										Clustering Type 2										Clustering Type 3										Clustering Type 4				
		c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c11	c1	c2	c3	c4
IE/Balto-Slavic	be	-	-	-	-	-	12.7	-	-	-	-	-	-	-	-	-	10.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	15.1
	bs	-	-	-	-	-	-	18.0	-	-	-	16.8	-	-	-	-	-	16.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	19.0	
	mk	-	-	-	-	-	21.4	-	-	-	-	20.0	-	-	-	-	-	16.9	17.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	18.3		
	it	-	-	-	-	-	22.4	-	-	-	-	21.1	-	-	-	-	-	22.4	21.1	24.0	-	-	-	-	-	-	-	-	-	-	-	-	-	22.8		
	cs	-	-	-	-	-	23.0	-	-	-	-	21.7	-	-	-	-	-	22.9	21.1	23.1	-	-	-	-	-	-	-	-	-	-	-	-	-	23.6		
	hr	-	-	-	-	-	27.5	-	-	-	-	27.5	-	-	-	-	-	25.6	22.9	23.5	-	-	-	-	-	-	-	-	-	-	-	-	-	28.3		
	sr	-	-	-	-	-	27.4	-	-	-	-	27.4	-	-	-	-	25.7	22.9	27.6	23.5	28.6	-	-	-	-	-	-	-	-	-	-	-	-	27.1		
	bg	-	-	-	-	-	19.4	-	-	-	-	18.3	-	-	-	-	-	20.8	20.8	22.1	-	-	-	-	-	-	-	-	-	-	-	-	-	30.4		
	pl	-	-	-	-	-	20.8	-	-	-	-	-	-	-	-	-	-	-	-	32.1	-	-	-	-	-	-	-	-	-	-	-	-	-	20.2		
	ru	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	21.4		
IE/Indo-Iranian	bn	-	-	9.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	ur	-	-	13.3	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	ku	-	-	9.4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	hi	-	-	13.2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
IE/Italic	fa	-	-	21.3	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	mr	-	-	8.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	gl	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	pt	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
IE/Indo-European	pt	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	ro	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	fr	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	es	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
IE/Indo-European	it	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	it	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	it	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	it	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		

Table 3: The results of clustering-based multilingual NMT with knowledge distillation for languages belong to IE/Balto-Slavic, IE/Indo-Iranian, and IE/Italic.

family	lang	Clustering Type 1										Clustering Type 2										Clustering Type 3										Clustering Type 4																			
		c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10										
Koreanic	ko	15.1	-	-	-	-	-	-	-	-	15.1	-	-	-	-	-	-	-	-	-	-	15.1	-	-	-	-	-	-	-	-	-	17.4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	ja	10.1	-	-	-	-	-	-	-	-	-	10.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	10.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Japonic	vi	-	21.6	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	21.6	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	21.3	-	-	-	-	-	-	-	-	-	20.3	-
	el	-	-	-	-	30.0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	31.3	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	29.6	-	-	-	-	-	-	-	-	-
IE/Hellenic	th	-	22.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	23.09	-
	sq	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	26.4	-	-	-	-	-	-	-	26.8	-

Table 5: The results of clustering-based multilingual NMT with knowledge distillation for languages belong to Koreanic, Japonic, Austroasiatic, IE/Hellenic, Kra-Dai, and IE/Albanian families.

Family	lang	Clustering Type 1										Clustering Type 2											Clustering Type 3											Clustering Type 4				
		c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c11	c1	c2	c3	c4												
IE/Armenian	hy	-	-	-	13.7	-	-	-	-	-	-	12.7	-	-	-	-	-	-	-	-	-	10.8	-	-	-	-	-	-	-	-	-	-	17.1	-	-	-	-	fa, ku, eu, hu, et, fi, hy, ka, ar, he, fr, cs, lt, de, nl, sv, da, nb, ru, be, pl, bg, sl, sr, uk, hr, mk, sk, ro, sq, el, es, gl, pt, it
		-	-	9.2	-	-	-	-	-	-	-	-	-	-	-	10.8	-	-	-	-	-	-	9.1	-	-	-	-	-	-	-	-	-	-	8.8	-	-	-	kk, az, tr, lt, ur, bn, mr, zh, my, ko, ja, mn, ta
Kartvelian	ka	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	vi, id, ms, th
		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	fa, ku, eu, hu, et, fi, hy, ka, ar, he, fr, cs, lt, de, nl, sv, da, nb, ru, be, pl, bg, sl, sr, uk, hr, mk, sk, ro, sq, el, es, gl, pt, it	
Mongolic	mm	5.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	eu, az, kk, mn, ur, mr, bn, ta, ku, bs, be, ms	
		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	zh, ja, ko, hu, tr	
Dravidian	ta	-	-	4.0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	nb, da, sv
		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	gl, fr, it, es, pt	
Isolate	ne	-	-	-	-	-	-	9.0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	eu, az, kk, mn, ur, mr, bn, ta, ku, bs, be, ms	
		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	hi, my, hy, ka	

Table 6: The results of clustering-based multilingual NMT with knowledge distillation for languages belong to IE/Armenian, Kartvelian, Mongolic, Dravidian, and Isolate families.

References

- Samuel Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [FAIRSEQ: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wj Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *ACL*.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. [When and why are pre-trained word embeddings useful for neural machine translation?](#) In *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, New Orleans, Louisiana. ACL.
- Fahimeh Saleh, Wray Buntine, and Gholamreza Hafari. 2020. [Collective wisdom: Improving low-resource neural machine translation using adaptive knowledge distillation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3413–3421.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with knowledge distillation](#). *arXiv preprint arXiv:1902.10461*.
- 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 30*, pages 5998–6008. Curran Associates, Inc.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020. [Balancing training for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8526–8537.