

Uncertainty-Aware Balancing for Multilingual and Multi-Domain Neural Machine Translation Training

Minghao Wu^{♡*} Yitong Li[♣] Meng Zhang[♣] Liangyou Li[♣]
Gholamreza Haffari[♡] Qun Liu[♣]

[♡]Monash University [♣]Huawei Technologies Co., Ltd. [♣]Huawei Noah's Ark Lab
{minghao.wu, gholamreza.haffari}@monash.edu
{liyitong3, zhangmeng92, liliangyou, qun.liu}@huawei.com

Abstract

Learning multilingual and multi-domain translation model is challenging as the heterogeneous and imbalanced data make the model converge inconsistently over different corpora in real world. One common practice is to adjust the share of each corpus in the training, so that the learning process is balanced and low-resource cases can benefit from the high-resource ones. However, automatic balancing methods usually depend on the intra- and inter-dataset characteristics, which is usually agnostic or requires human priors. In this work, we propose an approach, **MULTIUAT**, that dynamically adjusts the training data usage based on the model's uncertainty on a small set of trusted clean data for multi-corpus machine translation. We experiment with two classes of uncertainty measures on multilingual (16 languages with 4 settings) and multi-domain settings (4 for in-domain and 2 for out-of-domain on English-German translation) and demonstrate our approach **MULTIUAT** substantially outperforms its baselines, including both static and dynamic strategies. We analyze the cross-domain transfer and show the deficiency of static and similarity based methods.¹

1 Introduction

Text corpora are commonly collected from several different sources in different languages, raising the problem of learning a NLP system from the heterogeneous corpora, such as multilingual models (Wu and Dredze, 2019; Arivazhagan et al., 2019; Aharoni et al., 2019; Freitag and Firat, 2020; Arthur et al., 2021) and multi-domain models (Daumé III, 2007; Li et al., 2019; Deng et al., 2020; Jiang et al., 2020). A strong demand is to deploy a unified model for all the languages and domains, because

a unified model is much more resource-efficient, and knowledge learned from high-resource languages/domains (HRLs/HRDs) can be transferred to low-resource languages/domains (LRLs/LRDs).

One common issue on training models across corpora is that data from a variety of corpora are both heterogeneous (different corpora reveal different linguistic properties) and imbalance (the accessibility of training data varies across corpora). The standard practice to address this issue is to adjust the training data distribution heuristically by up-sampling the training data from LRLs/LRDs (Arivazhagan et al., 2019; Conneau et al., 2020).

Arivazhagan et al. (2019) rescale the training data distribution with a heuristic temperature term and demonstrate that the ideal temperature can substantially improve the overall performance. However, the optimal value for such heuristics is both hard to find and varies from one experimental setting to another (Wang and Neubig, 2019; Wang et al., 2020a,b). Wang et al. (2020a) and Wang et al. (2020b) hypothesize that the training data instances that are similar to the validation set can be more beneficial to the evaluation performance and propose a general reinforcement-learning framework *Differentiable Data Selection* (DDS) that automatically adjusts the importance of data points, whose reward is the cosine similarity of the gradients between a small set of trusted clean data and training data. They instantiate this framework on multilingual NMT, known as MULTIDDS, to dynamically weigh the importance of language pairs. Both the hypothesis and the proposed approach rely on the assumption that knowledge learned from one corpus can always be beneficial to the other corpora.

However, their assumption does not always hold. If the knowledge learned from one corpus is not able to be transferred easily or is useless to the other corpora, this approach fails. Unlike cosine similarity, *model uncertainty* is free from the aforementioned assumption on cross-corpus transfer. From

*Work done during the internship at Huawei Noah's Ark Lab.

¹Code available at <https://github.com/huawei-noah/noah-research/tree/master/noahnmt/multiuat>

a Bayesian viewpoint, the model parameters can be considered as a random variable that describes the dataset. If one dataset is well-described by the model parameters, its corresponding model uncertainty is low, and vice versa. This nature makes the model uncertainty an ideal option to weigh the datasets.

In this work, we propose an approach MULTI-UAT that leverages the model uncertainty as the reward to dynamically adjust the sampling probability distribution over multiple corpora. We consider the model parameter as a random variable that describes the multiple training corpora. If one corpus is well-described by the model compared with other corpora, we spare more training efforts to the other poorly described corpora. We conduct extensive experiments on multilingual NMT (16 languages with 4 settings) and multi-domain NMT (4 for in-domain and 2 for out-of-domain), comparing our approach with heuristic static strategy and dynamic strategy. In multilingual NMT, we improve the overall performance from +0.83 BLEU score to +1.52 BLEU score among 4 settings, comparing with the best baseline. In multi-domain NMT, our approach improves the in-domain overall performance by +0.58 BLEU score comparing with the best baseline and achieves the second best out-of-domain overall performance. We also empirically illustrate the vulnerability of cosine similarity as the reward in the training among multiple corpora.

2 Preliminaries

Standard NMT A standard NMT model, parameterized by θ , is commonly trained on one language pair $D_{\text{trn}}^o = \{(\mathbf{x}, \mathbf{y})_i\}_{i=1}^M$ from one domain. The objective is to minimize the negative log-likelihood of the training data with respect to θ :

$$\mathcal{L}_s(D_{\text{trn}}^o; \theta) = - \sum_{i=1}^M \log p(\mathbf{y} | \mathbf{x}; \theta). \quad (1)$$

Multi-corpus NMT Both multilingual NMT and multi-domain NMT can be summarized as multi-corpus NMT that aims to build a unified translation system to maximize the overall performance across all the language pairs or domains. Formally, let us assume we are given a number of datasets $D_{\text{trn}} = \{D_{\text{trn}}^j\}_{j=1}^N$ from N languages pairs or domains, in which $D_{\text{trn}}^j = \{(\mathbf{x}, \mathbf{y})_i^j\}_{i=1}^{M_j}$, where M_j is size of j -th language/domain. Similar to Equation 1, a simple way of training multi-corpus NMT model

is to treat all instances equally:

$$\mathcal{L}(D_{\text{trn}}; \theta) = \sum_{j=1}^N \mathcal{L}_s(D_{\text{trn}}^j; \theta). \quad (2)$$

Heuristic strategy for multi-corpus training

In practice, Equation 2 can be reviewed as training using mini-batch sampling according to the proportion of these corpora, $q(n) = \frac{M_n}{\sum_{i=1}^N M_i}$, and thus we minimize:

$$\mathcal{L}(D_{\text{trn}}; \theta, q(n)) = \mathbb{E}_{n \sim q(n)} [\mathcal{L}_s(D_{\text{trn}}^n; \theta)]. \quad (3)$$

However, this simple training method does not work well in real cases, where low-resource tasks are under-trained.

A heuristic static strategy is to adjust the proportion exponentiated by a temperature term τ (Ari-vazhagan et al., 2019; Conneau et al., 2020):

$$q_\tau(n) = \frac{q(n)^{1/\tau}}{\sum_{i=1}^N q(i)^{1/\tau}}. \quad (4)$$

And the loss function for multi-corpus training can be re-formulated as:

$$\mathcal{L}(D_{\text{trn}}; \theta, q_\tau(n)) = \mathbb{E}_{n \sim q_\tau(n)} [\mathcal{L}_s(D_{\text{trn}}^n; \theta)]. \quad (5)$$

Specifically, $\tau = 1$ or $\tau = \infty$ is equivalent to proportional (Equation 2) or uniform sampling respectively.

Differentiable Data Selection (DDS) Wang et al. (2020a)

Wang et al. (2020a) propose a general framework that automatically weighs training instances to improve the performance while relying on an independent set of held-out data D_{dev} . Their framework consists of two major components, the model θ and the scorer network ψ . The scorer network ψ is trained to assign a sampling probability to each training instance, denoted as $p_\psi(\mathbf{x}, \mathbf{y})$, based on its contribution to the validation performance. The training instance that contributes more to the validation performance is assigned a higher probability and more likely to be used for updating the model θ . This strategy aims to maximize the overall performance over D_{dev} and is expected to generalize well on the unseen D_{tst} with the assumption of the independence and identical distribution between D_{dev} and D_{tst} . Therefore, the objective can be formulated as:

$$\begin{aligned} \psi &= \underset{\psi}{\operatorname{argmin}} \mathcal{L}(D_{\text{dev}}; \theta(\psi)) \\ \theta(\psi) &= \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p_\psi(\mathbf{x}, \mathbf{y})} [\mathcal{L}(D_{\text{trn}}; \theta)], \end{aligned} \quad (6)$$

and ψ and θ are updated iteratively using *bilevel optimization* (Colson et al., 2007; von Stackelberg et al., 2011).

3 Methodology

In this work, we leverage the idea of DDS under the multi-corpus scenarios. We utilize a differentiable domain/language scorer ψ to weigh the training corpora. To learn ψ , we exploit the model uncertainty to measure the model’s ability over the target corpus. Below, we elaborate on the details of our method.

3.1 Model Uncertainty

Model uncertainty can be a measure that indicates whether the model parameters θ are able to describe the data distribution well (Kendall and Gal, 2017; Dong et al., 2018; Xiao and Wang, 2019). Bayesian neural networks can be used for quantifying the model uncertainty (Buntine and Weigend, 1991), which models the θ as a probabilistic distribution with constant input and output.

From the Bayesian point of view, θ is interpreted as a random variable with the prior $p(\theta)$. Given a dataset D , the posterior $p(\theta|D)$ can be obtained via Bayes’ rule. However, the exact Bayesian inference is intractable for neural networks, so that it is common to place the approximation $q(\theta)$ to the true posterior $p(\theta|D)$. Several variational inference methods have been proposed (Graves, 2011; Blundell et al., 2015; Gal and Ghahramani, 2016).

In this work, we leverage Monte Carlo Dropout (Gal and Ghahramani, 2016) to obtain samples of sentence-level translation probability. To quantify the model uncertainty when the model makes predictions, we treat the sentence-level translation probability as random variable. We run K forward passes with a random subset of model parameters θ deactivated, which is equivalent to drawing samples from the random variable, and average the samples as the estimate of the model uncertainty.² Consider an ensemble of models $\{p_{\theta_k}(\mathbf{y}|\mathbf{x})\}_{k=1}^K$ sampled from the approximate posterior $q(\theta)$, the predictive posterior can be obtained by taking the expectation over multiple inferences:

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}, D) &\approx \mathbb{E}_{\theta \sim q(\theta)} [p(\mathbf{y}|\mathbf{x}, \theta)] \\ &\approx \frac{1}{K} \sum_{k=1}^K p_{\theta_k}(\mathbf{y}|\mathbf{x}). \end{aligned} \quad (7)$$

² K is set to 30 in our work.

Algorithm 1: Training with MULTIUAT

Input : $D^n = \{(\mathbf{x}_m^n, \mathbf{y}_m^n)\}_{m=1}^{M_n}$, N corpora with the size of M_n for the n -th corpus; S , update frequency of ψ ; J , total training steps;

Output : The converged model θ

```

1 Initialize  $p_\psi(n; D_{\text{trn}})$  as Equation 4 with  $\tau = 1$ ;
2 for  $i=0$  to  $J$  do
3    $\tilde{n} \sim p_\psi(n)$ ;
4   sample batch  $(\mathbf{x}, \mathbf{y}) \sim D_{\text{trn}}^{\tilde{n}}$ ;
5    $\theta \leftarrow \theta - \alpha \cdot \nabla_{\theta} \mathcal{L}(\mathbf{y}|\mathbf{x}; \theta)$ ;
6   if  $i \% S == 0$  then
7     for  $n=1$  to  $N$  do
8        $(\mathbf{x}', \mathbf{y}') \sim D_{\text{dev}}^n$ ;
9       Compute reward/uncertainty  $R(n)$  for  $D_{\text{dev}}^n$  as in Section 3.3;
10    end
11     $\psi \leftarrow \psi - \sum_{n=1}^N R(n) \cdot \nabla_{\psi} \log p_\psi(n)$ ;
12  end
13 end

```

3.2 Uncertainty-Aware Training

To make the training more efficient and stable, MULTIUAT leverages the scorer network ψ to dynamically adjust the sampling probability distribution of domains/languages.³

We present the pseudo-code for training with MULTIUAT in Algorithm 1. MULTIUAT firstly parameterizes the initial sampling probability distribution for multi-corpus training with ψ as Equation 4 with the warm-up temperature $\tau = 1$. For the computational efficiency, the scorer network ψ is updated for every S steps. When updating ψ , we randomly draw one mini-batch from each validation set $\{D_{\text{dev}}^i\}_{i=1}^N$ and compute the corresponding uncertainty measure as in Section 3.3 with Monte Carlo Dropout to approximate the model uncertainty towards this corpus, assuming the validation set is representative enough for its corresponding true distribution. The corpus associated with high uncertainty is considered to be relatively poorly described by the model θ and its sampling probability will be increased. The model θ is updated by mini-batch gradient descent between two updates

³We parameterize the scorer ψ following Wang et al. (2020b).

of ψ , like common gradient-based optimization, and hence the objective is formulated as follows:

$$\begin{aligned} \psi &= \underset{\psi}{\operatorname{argmin}} \mathcal{L}(D_{\text{dev}}; \theta(\psi)) \\ \theta(\psi) &= \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{n \sim p_{\psi}(n)} [\mathcal{L}(D_{\text{tm}}^n; \theta)]. \end{aligned} \quad (8)$$

A considerable problem here is Equation 6 is not directly differentiable w.r.t. the scorer ψ . To tackle this problem, reinforcement learning (RL) with suitable reward functions is required (Fang et al., 2017; Wang et al., 2020a):

$$\psi \leftarrow \psi - \sum_{n=1}^N R(n) \cdot \nabla_{\psi} \log p_{\psi}(n). \quad (9)$$

Details for the reward functions $R(n)$ are depicted at Section 3.3 and the update of ψ follows the REINFORCE algorithm (Williams, 1992).

3.3 Uncertainty Measures

We explore the utility of two groups of model uncertainty measures: *probability-based* and *entropy-based* measures at the sentence level (Wang et al., 2019; Fomicheva et al., 2020; Malinin and Gales, 2021).

Probability-Based Measures We explore four probability-based uncertainty measures following the definition of Wang et al. (2019). For the sampled model parameters θ_k , with the teacher-forced decoding, we note the predicted probability of the t -th position as:

$$\hat{y}_{n,t} = \underset{y}{\operatorname{argmax}} p(y | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k), \quad (10)$$

where we have used the ground truth prefix $\mathbf{y}_{n,<t}$ in the conditioning context. We then define the reward function as the following uncertainty measures:

- *Predicted Translation Probability* (PRETP): The predicted probability of the sentence,

$$R_{\text{PRETP}}(n; \theta_k) = 1 - \prod_{t=1}^T p(\hat{y}_{n,t} | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k).$$

- *Expected Translation Probability* (EXPTP): The expectation of the distribution of maximal position-wise translation probability,

$$R_{\text{EXPTP}}(n; \theta_k) = 1 - \mathbb{E} [p(\hat{y}_{n,t} | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k)].$$

- *Variance of Translation Probability* (VARTP): The variance of the distribution of maximal position-wise translation probability,

$$R_{\text{VARTP}}(n; \theta_k) = \operatorname{Var}[p(\hat{y}_{n,t} | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k)].$$

- *Combination of Expectation and Variance* (COMEV):

$$R_{\text{COMEV}}(n; \theta_k) = \frac{\operatorname{Var}[p(\hat{y}_{n,t} | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k)]}{\mathbb{E}[p(\hat{y}_{n,t} | \mathbf{x}_n, \mathbf{y}_{n,<t}; \theta_k)]}.$$

Entropy-Based Measures Malinin and Gales (2021) consider the uncertainty estimation for auto-regressive models at the token-level and sequence-level and treat the entropy of the posterior as the total uncertainty in the prediction of \mathbf{y} . Following their interpretation, we leverage the *entropy* as the measure of the model uncertainty.

Drawing a pair of sentence (\mathbf{x}, \mathbf{y}) with T target tokens from the n -th corpus D^n , the reward function is defined as the averaged entropy over all the positions:

$$R(n; \theta) = \frac{1}{T} \sum_{t=1}^T \sum_{v=1}^{\mathcal{V}} p(y_{n,t,v}) \log p(y_{n,t,v}). \quad (11)$$

where \mathcal{V} is the vocabulary size and $p(y_{n,t,v})$ stands for the predicted conditional probability $p(y_{n,t,v} | \mathbf{x}, \mathbf{y}_{n,<t,v}; \theta_k)$ on the v -th word in the vocabulary.

In this work, we explore the utility of two entropy-based uncertainty measures as follows:

- *Entropy of the sentence* (ENTSENT): The average entropy of the sentence as defined in Equation 11.
- *Entropy of EOS* (ENTEOS): The entropy of the symbol EOS in the sentence as defined in Equation 11 where $t = T$.

Following Equation 7, we have the final reward by multiple sampled θ_k for each uncertainty reward respectively:

$$R(n) = \frac{1}{K} \sum_{k=1}^K R_{\{\cdot\}}(n; \theta_k). \quad (12)$$

4 Experimental Setup

4.1 Baselines

We compare MULTIUAT with both static and dynamic strategies as follows:

Heuristics We run experiments with proportional (PROP., $\tau = 1$), temperature (TEMP., $\tau = 5$) and uniform (UNI., $\tau = \infty$) in Equation 4 following Wang et al. (2020b).

MULTIDDS-S We compare with the best model MULTIDDS-S proposed by Wang et al. (2020b) over multilingual NMT tasks. Its reward for the n -th corpus is defined using cosine similarity:

$$R_{\cos}(n) = \frac{1}{N} \sum_{i=1}^N \cos(\nabla_{\theta} \mathcal{L}(D_{\text{dev}}^i), \nabla_{\theta} \mathcal{L}(D_{\text{trn}}^n)). \quad (13)$$

4.2 Multilingual Setup

We follow the identical setup as Wang et al. (2020b) in the multilingual NMT. The model is trained on two sets of language pairs based on the language diversity.

Related 4 LRLs (Azerbaijani: `aze`, Belarusian: `bel`, Glacian: `glg`, Slovak: `slk`) and a related HRL for each LRL (Turkish: `tur`, Russian: `rus`, Portuguese: `por`, Czech: `ces`).

Diverse 8 languages with varying amounts of data, picked without consideration for relatedness (Bosnian: `bos`, Marathi: `mar`, Hindi: `hin`, Macedonian: `mkd`, Greek: `ell`, Bulgarian: `bul`, French: `fra`, Korean: `kor`).

We run many-to-one (M2O, translating 8 languages to English) and one-to-many (O2M, translating English to 8 languages) translations for both diverse and related setups.⁴

4.3 Multi-Domain Setup

We run experiments on English-German translation and collect six corpora from WMT2014 (Bojar et al., 2014) and the Open Parallel Corpus (Tiedemann, 2012), 4 for in-domain and 2 for out-of-domain:

In-Domain (ID) (i) WMT, from WMT2014 translation task (Bojar et al., 2014) with the concatenation from `newstest2010` to `newstest2013` for validation and `newstest2014` for testing; (ii) `Tanzil`,⁵ a collection of Quran translations; (iii) `EMEA`,⁶ a parallel corpus from the European Medicines Agency; (iv) `KDE`,⁷ a parallel corpus of KDE4 localization files.

⁴Refer to Wang et al. (2020b) for dataset statistics.

⁵<https://opus.nlpl.eu/Tanzil.php>

⁶<https://opus.nlpl.eu/EMEA.php>

⁷<https://opus.nlpl.eu/KDE4.php>

		Train	Valid	Test
ID	WMT	3,950K	11K	3K
	Tanzil	449K	3K	3K
	EMEA	277K	3K	3K
	KDE	135K	3K	3K
OOD	QED	-	-	3K
	TED	-	-	3K

Table 1: Dataset statistics of multi-domain corpora.

Out-Of-Domain (OOD) (i) `QED`,⁸ a collection of subtitles for educational videos and lectures (Abdelali et al., 2014); (ii) `TED`,⁹ a parallel corpus of TED talk subtitles. These two domains are only used for out-of-domain evaluation.

All these corpora are first tokenized by Moses (Koehn et al., 2007) and processed into sub-word units by BPE (Sennrich et al., 2016) with 32K merge operations. Sentence pairs that are duplicated and violates source-target ratio of 1.5 are removed. The validation sets and test sets are randomly sampled, except for WMT. The dataset statistics are listed in Table 1.

4.4 Model Architecture

We believe all the approaches involved in this work, including the baseline approaches and MULTIUAT, are model-agnostic. To validate this idea, we experiment two variants of transformer (Vaswani et al., 2017). For multilingual NMT, the model architecture is a transformer with 4 attention heads and 6 layers.¹⁰ And for multi-domain NMT models, we use the standard transformer-base with 8 attention heads and 6 layers.¹¹ All the models in this work are implemented by fairseq (Ott et al., 2019).

4.5 Evaluation

We report detokenized BLEU (Papineni et al., 2002) using SacreBLEU (Post, 2018) with statistical significance given by Koehn (2004).¹² μ_{BLEU} is the macro average of BLEU scores within the same setting, with the assumption that all the language pairs/domains are equally important.

⁸<https://opus.nlpl.eu/QED.php>

⁹<https://opus.nlpl.eu/TED2013.php>

¹⁰Signature:

`multilingual_transformer_iwslt_de_en`

¹¹Signature: `transformer`

¹²Signature: `BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.14`

	Average	Multilingual				Multi-Domain	
		Related		Diverse		ID	OOD
		M2O	O2M	M2O	O2M		
PROP. ($\tau = 1$)	22.58	24.88	15.49	26.68	16.79	35.69	25.15
TEMP. ($\tau = 5$)	22.80	24.00	16.61	26.01	17.94	36.92	23.46
UNI. ($\tau = \infty$)	21.86	22.63	15.54	24.81	16.86	36.85	22.22
MULTIDDS-S	23.58	25.52	17.32	27.00	18.24	36.42	22.74
MULTIUAT							
+ PRETP	24.57†	26.30†	18.36†	<u>27.82</u> †	19.57†	37.50 †	23.77
+ EXPPTP	<u>24.62</u> †	<u>26.36</u> †	18.64 †	27.74†	19.59†	37.29	<u>23.89</u>
+ VARTP	24.44†	26.39 †	18.58†	27.83 †	19.62†	35.49	23.75
+ COMEV	24.57†	26.26†	18.58†	27.75†	19.67†	37.25	23.36
+ ENTSENT	<u>24.62</u> †	26.34†	<u>18.63</u> †	27.71†	<u>19.68</u> †	37.27	23.86
+ ENTEOS	24.65 †	26.28†	18.56†	<u>27.82</u> †	19.76 †	<u>37.44</u> †	23.80

Table 2: μ_{BLEU} for the settings on multilingual and multi-domain NMT. ‘‘Average’’ is the macro average of all the BLEU scores for both multilingual and multi-domain settings at Appendix A. Best results are highlighted in **bold** and second best results are highlighted in underline. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ using paired Student’s t-test.

5 Main Results

The summarized results for both multilingual and multi-domain NMT are presented in Table 2.¹³ The complete results with statistical significance can be found in Appendix A.

Multilingual NMT Overall, dynamic strategies (MULTIDDS-S and MULTIUAT) demonstrate their superiority against heuristic static strategies. As shown in Table 2, the optimal τ of heuristic static strategies varies as the combination of corpora changes. For example, proportional sampling yields best performance on M2O settings, yet achieves the worst performance on O2M settings among heuristic static strategies. Dynamic strategies are free from adjusting the data usage by tuning the τ . MULTIDDS-S marginally outperforms heuristic static strategies. MULTIUAT with various uncertainty measures reaches the best performance in all four settings. Based on the detailed results in Appendix A, we can observe that MULTIUAT appears to be more favorable to HRLs.

Multi-domain NMT MULTIUAT outperforms all its baselines on in-domain evaluation and achieves the second best performance on out-of-domain evaluation. MULTIUAT with PRETP achieves the optimal balance on in-domain evaluation and the one with EXPPTP achieves the second best performance on out-of-domain evaluation. However, MULTIDDS-S performs poorly on multi-

domain NMT and is even outperformed by some heuristic static strategies.

Based on the detailed results in Appendix A, we can observe that the higher sampling probability for certain domain is commonly but not always positively correlated to the corresponding in-domain performance. Uniformly sampling minibatches from domains does not result in the best performance on LRDs, because the LRDs with too much up-sampling are not able to fully leverage the knowledge from the HRDs.

6 Analysis

Wang et al. (2020b) conduct exhaustive analyses on multilingual NMT and most of our observations are consistent with theirs.¹⁴ Hence, we focus more on analyzing the results on multi-domain NMT.

6.1 Comparison of Uncertainty Measures

We explore the utility of different uncertainty measures and display the summarized results in Table 2. Different uncertainty measures deliver different results. We do not observe one uncertainty measure that consistently outperforms others. The probability-based uncertainty measures seem to be more sensitive to the intra- and inter-dataset characteristics, and perform well on either multilingual NMT or multi-domain NMT. MULTIUAT with the uncertainty measure of VARTP performs substantially worse than other uncertainty measures in

¹³We report the multilingual results of the heuristic approach and MULTIDDS-S from Wang et al. (2020b) in Table 2.

¹⁴The multilingual results of MULTIDDS-S used for analysis are provided by our own implementation with the hyperparameters provided by Wang et al. (2020b).

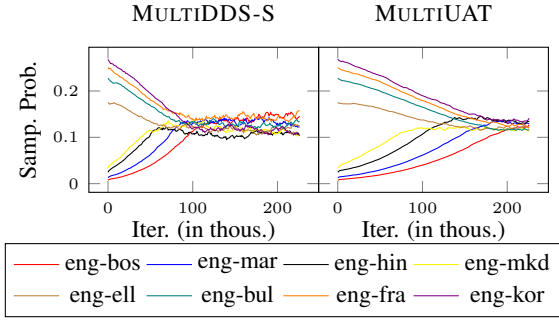


Figure 1: Iteration (Iter.)-sampling probability of MULTIDDS-S and MULTIUAT under the multilingual O2M-diverse setting.

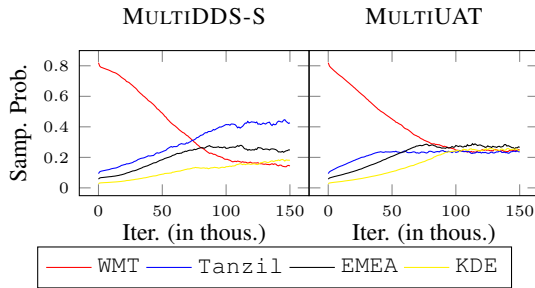


Figure 2: Iteration (Iter.)-sampling probability of MULTIDDS-S and MULTIUAT over multi-domain NMT training sets.

multi-domain NMT. In contrast to the probability-based uncertainty measures, the entropy-based uncertainty measures are more robust to the change of datasets and deliver relatively stable improvements. We also find out that MULTIUAT with the uncertainty measures demonstrate better out-of-domain generalization in the multi-domain NMT, compared with its baselines.

Based on the detailed results in Appendix A, MULTIUAT with the entropy-based uncertainty measures demonstrates better robustness against the change of datasets. Therefore, we mainly compare MULTIUAT with the uncertainty measure of ENTEOS against the baselines in the following analyses, based on the macro-average results on both multilingual and multi-domain NMT.

6.2 Learned Distribution for Language Pairs/Domains

We visualize the change of sampling distribution, w.r.t. the training iterations, of the multilingual O2M-diverse (Figure 1) and multi-domain (Figure 2) setting. In both figures, MULTIDDS-S and MULTIUAT gradually increase the usage of LRLs/LRDs and decrease the usage of

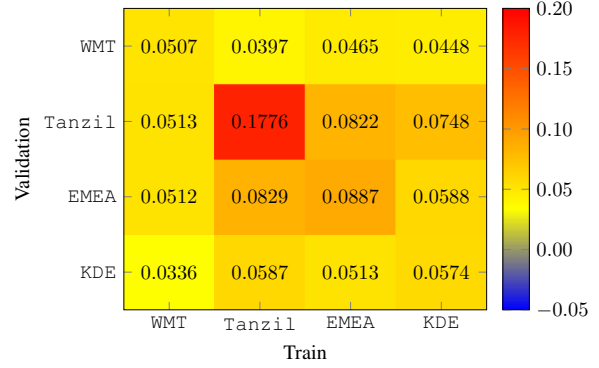


Figure 3: Averaged cosine similarity matrix between train and validation mini-batch during the optimization of MULTIDDS-S.

HRLs/HRDs. In the multilingual NMT, we observe that the learned distributions by both MULTIDDS-S and MULTIUAT converge from proportional sampling to uniform sampling with a mild trend to divergence in the one given by MULTIDDS-S. In the multi-domain NMT, MULTIUAT illustrates the consistent adjustment as the trend illustrated in multilingual O2M-diverse setting, but the learned distribution given by MULTIDDS-S is overwhelmed by Tanzil.

The model uncertainty focuses on how well the dataset is described by the model θ , instead of the interference among datasets, so that MULTIUAT is free from the assumption on the cross-corpus transference and not affected by Tanzil.

6.3 Why Cosine Similarity Fails?¹⁵

A natural question is raised after seeing Figure 2: why does Tanzil overwhelm the sampling distribution by MULTIDDS-S in multi-domain NMT?

As in Equation 13, MULTIDDS-S computes pairwise cosine similarities for all the language pairs/domains using sampled mini-batches between D_{trn} and D_{dev} to update the sampling probability. We average all the cosine similarity matrices during the training and visualize the averaged matrix in Figure 3. As visualized, Tanzil is a highly self-correlated domain whose cosine similarity is about at least two times larger than the other values in the matrix. This leads to a very high reward on Tanzil, and the sampling probability of Tanzil in MULTIDDS-S keeps increasing to more than 40% in Figure 2.

¹⁵In our preliminary study, we investigate various similarity-based rewards, such as dot product, but we do not observe significant difference.

	τ	ID					OOD		
		μ_{BLEU}	WMT	Tanzil	EMEA	KDE	μ_{BLEU}	QED	TED
MULTIDDS-S	1	36.42	21.76	37.41	51.92	34.60	22.74	20.30	25.18
	5	36.16	19.50	37.55	52.04	35.56	21.60	19.59	23.61
	∞	36.03	18.90	37.45	51.91	<u>35.85</u>	21.47	19.37	23.57
MULTIUAT	1	37.44	22.14	38.38†	52.81 †	36.44 †	23.80	21.34 †	26.26 †
	5	<u>37.04</u>	20.34	<u>39.40</u> †	<u>52.68</u>	35.73	22.46	19.97	24.95
	∞	36.93	20.09	39.84 †	52.03	35.74	22.14	19.79	24.48

Table 3: Effects of sampling priors (τ) for MULTIDDS-S and MULTIUAT in multi-domain NMT. Best results are highlighted in **bold** and second best results are highlighted in underline. † indicates the improvement for the corresponding result against best MULTIDDS-S result is statistically significant at $p < 0.05$ given by Koehn (2004).

	WMT	Tanzil	EMEA	KDE
WMT	25.09	12.76	23.60	24.95
Tanzil	0.22	40.66	0.07	0.14
EMEA	4.78	1.39	54.25	7.01
KDE	4.19	1.74	9.54	30.71

Table 4: The cross-domain evaluation for the single-domain NMT models. BLEU scores in each row are produced by the corresponding single-domain NMT model. Best results are highlighted in **bold**.

However, is Tanzil highly beneficial to the overall performance? To probe the cross-domain generalization, we train four single-domain NMT models on each in-domain corpus and evaluate these models on all the in-domain test sets, and the results are presented in Table 4. We can observe that the knowledge learned from WMT can be generalized to other domains, but the knowledge learned from Tanzil is almost not beneficial to other domains.

Therefore, MULTIDDS-S with the data-dependent cosine similarity reward is vulnerable to the change of datasets and can be possibly overwhelmed by a special dataset like Tanzil, since the cross-corpus transfer is intractable.

6.4 Effects of Sampling Priors

Both MULTIDDS-S and MULTIUAT initialize the sampling probability distribution to proportional distribution (line 1 in Algorithm 1). We investigate how the prior sampling distribution affects the performance and present the results in Table 3. We can observe that the prior sampling distribution can affect the overall performance. For both MULTIDDS-S and MULTIUAT, the overall results on both in-domain and out-of-domain evaluation are negatively correlated with the prior τ .

We also visualize the change of sampling probability of KDE given by MULTIDDS-S and MUL-

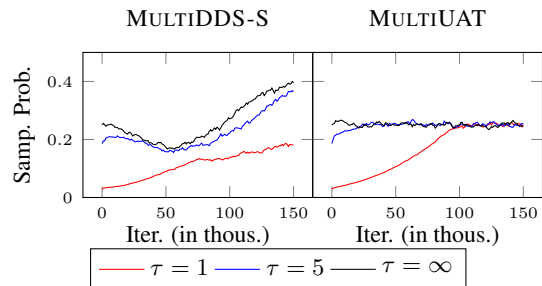


Figure 4: Iteration (Iter.)-sampling probability over KDE of MULTIDDS-S and MULTIUAT with different temperature priors.

MULTIUAT with different prior sampling distributions in Figure 4. The learned sampling distribution by MULTIUAT always converges to uniform distribution, regardless of the change of prior sampling distribution. However, the change of priors significantly affects the learned sampling distribution of MULTIDDS-S.

7 Related Work

Multi-corpus NLP Multilingual training has been particularly prominent in recent advances driven by the demand of training a unified model for all the languages (Dong et al., 2015; Plank et al., 2016; Johnson et al., 2017; Arivazhagan et al., 2019). Freitag and Firat (2020) extend current English-centric training to a many-to-many setup without sacrificing the performance on English-centric language pairs. Wang et al. (2021) improve the multilingual training by adjusting gradient directions based on gradient similarity. Existing works on multi-domain training commonly attempt to leverage architectural domain-specific components or auxiliary loss (Sajjad et al., 2017; Tars and Fishel, 2018; Zeng et al., 2018; Li et al., 2018; Deng et al., 2020; Jiang et al., 2020). These

approaches commonly do not explore much on the training proportion across domains and are limited to in-domain prediction and less generalizable to unseen domains. Zaremoodi and Haffari (2019) dynamically balance the importance of tasks in multi-task NMT to improve the low-resource NMT performance. Vu et al. (2021) leverage a pre-trained language model to select useful monolingual data from either source language or target language to perform unsupervised domain adaptation for NMT models. Our work is directly related to Wang et al. (2020a) and Wang et al. (2020b) that leverage cosine similarity of gradients as a reward to dynamically adjust the data usage in the multilingual training.

Model uncertainty Estimating the sequence and word-level uncertainty via Monte Carlo Dropout (Gal and Ghahramani, 2016) has been investigated for NMT (Xiao et al., 2020; Wang et al., 2019; Fomicheva et al., 2020; Malinin and Gales, 2021). Wang et al. (2019) exploit model uncertainty on back-translation to reduce the noise in the back-translated corpus. Xiao et al. (2020) and Malinin and Gales (2021) investigate to leverage model uncertainty to detect out-of-distribution translations. Fomicheva et al. (2020) summarize several measures to estimate quality of translated sentences, including the model uncertainty. Our work exploits the uncertainty measures as suggested by Wang et al. (2019) and Malinin and Gales (2021).

8 Conclusion

In this work, we propose MULTIUAT, a general model-agnostic framework that learns to automatically balance the data usage to achieve better overall performance on multiple corpora based on model uncertainty. We run extensive experiments on both multilingual and multi-domain NMT, and empirically demonstrate the effectiveness of our approach. Our approach substantially outperforms other baseline approaches. We empirically point out the vulnerability of a comparable approach MULTIDDS-S (Wang et al., 2020b).

We focus on the problem that dynamically balances text corpora collected from heterogeneous sources in this paper. However, the heterogeneity of text corpora is far beyond the languages and domains which are discussed in this work. For example, the quality of datasets is not covered. We leave the study on the quality of datasets to the future work.

Acknowledgments

We would like to thank the anonymous reviewers for providing constructive and helpful comments on this work. This work is partly supported by the ARC Future Fellowship FT190100039.

References

- Ahmed Abdelali, Francisco Guzman, Hassan Sajjad, and Stephan Vogel. 2014. [The AMARA corpus: Building parallel language resources for the educational domain](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1856–1862, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). 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 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Philip Arthur, Dongwon Ryu, and Gholamreza Haffari. 2021. [Multilingual simultaneous neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4758–4766, Online. Association for Computational Linguistics.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. [Weight uncertainty in neural network](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622, Lille, France. PMLR.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Wray L. Buntine and Andreas S. Weigend. 1991. [Bayesian back-propagation](#). *Complex Syst.*, 5(6).

- Benoît Colson, Patrice Marcotte, and Gilles Savard. 2007. [An overview of bilevel optimization](#). *Ann. Oper. Res.*, 153(1):235–256.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Hal Daumé III. 2007. [Frustratingly easy domain adaptation](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 256–263, Prague, Czech Republic. Association for Computational Linguistics.
- Yongchao Deng, Hongfei Yu, Heng Yu, Xiangyu Duan, and Weihua Luo. 2020. [Factorized transformer for multi-domain neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 4221–4230. Association for Computational Linguistics.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1723–1732, Beijing, China. Association for Computational Linguistics.
- Li Dong, Chris Quirk, and Mirella Lapata. 2018. [Confidence modeling for neural semantic parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 743–753, Melbourne, Australia. Association for Computational Linguistics.
- Meng Fang, Yuan Li, and Trevor Cohn. 2017. Learning how to active learn: A deep reinforcement learning approach. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 595–605. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. [Unsupervised quality estimation for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Markus Freitag and Orhan Firat. 2020. [Complete multilingual neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 550–560, Online. Association for Computational Linguistics.
- Yarin Gal and Zoubin Ghahramani. 2016. [Dropout as a bayesian approximation: Representing model uncertainty in deep learning](#). In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059. JMLR.org.
- Alex Graves. 2011. [Practical variational inference for neural networks](#). In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, pages 2348–2356.
- Wenbin Jiang, Mengfei Guo, Yufeng Chen, Ying Li, Jinan Xu, Yajuan Lyu, and Yong Zhu. 2020. [Multi-view classification model for knowledge graph completion](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 726–734, Suzhou, China. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Alex Kendall and Yarin Gal. 2017. [What uncertainties do we need in bayesian deep learning for computer vision?](#) In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5574–5584.
- 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*.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.

- Yitong Li, Timothy Baldwin, and Trevor Cohn. 2018. [What’s in a domain? learning domain-robust text representations using adversarial training](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 474–479, New Orleans, Louisiana. Association for Computational Linguistics.
- Yitong Li, Timothy Baldwin, and Trevor Cohn. 2019. [Semi-supervised stochastic multi-domain learning using variational inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1923–1934, Florence, Italy. Association for Computational Linguistics.
- Andrey Malinin and Mark Gales. 2021. [Uncertainty estimation in autoregressive structured prediction](#). In *International Conference on Learning Representations*.
- 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. [Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 412–418, Berlin, Germany. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Hassan Sajjad, Nadir Durrani, Fahim Dalvi, Yonatan Belinkov, and Stephan Vogel. 2017. [Neural machine translation training in a multi-domain scenario](#). *CoRR*, abs/1708.08712.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [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.
- Sander Tars and Mark Fishel. 2018. [Multi-domain neural machine translation](#). *CoRR*, abs/1805.02282.
- 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).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Heinrich von Stackelberg et al. 2011. [Market structure and equilibrium](#). *Springer Books*.
- Thuy-Trang Vu, Xuanli He, Dinh Phung, and Gholamreza Haffari. 2021. [Generalised unsupervised domain adaptation of neural machine translation with cross-lingual data selection](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- Shuo Wang, Yang Liu, Chao Wang, Huanbo Luan, and Maosong Sun. 2019. [Improving back-translation with uncertainty-based confidence estimation](#). 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 791–802, Hong Kong, China. Association for Computational Linguistics.
- Xinyi Wang and Graham Neubig. 2019. [Target conditioned sampling: Optimizing data selection for multilingual neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5823–5828, Florence, Italy. Association for Computational Linguistics.
- Xinyi Wang, Hieu Pham, Paul Michel, Antonios Anastopoulos, Jaime G. Carbonell, and Graham Neubig. 2020a. [Optimizing data usage via differentiable rewards](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9983–9995. PMLR.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020b. [Balancing training for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8526–8537, Online. Association for Computational Linguistics.
- Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. 2021. [Gradient vaccine: Investigating and im-](#)

proving multi-task optimization in massively multilingual models. In *International Conference on Learning Representations*.

Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8(3–4):229–256.

Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT. 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 833–844, Hong Kong, China. Association for Computational Linguistics.

Tim Z. Xiao, Aidan N. Gomez, and Yarin Gal. 2020. Wat zei je? detecting out-of-distribution translations with variational transformers. *CoRR*, abs/2006.08344.

Yijun Xiao and William Yang Wang. 2019. Quantifying uncertainties in natural language processing tasks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 7322–7329. AAAI Press.

Poorya Zareemoodi and Gholamreza Haffari. 2019. Adaptively scheduled multitask learning: The case of low-resource neural machine translation. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 177–186, Hong Kong. Association for Computational Linguistics.

Jiali Zeng, Jinsong Su, Huating Wen, Yang Liu, Jun Xie, Yongjing Yin, and Jianqiang Zhao. 2018. Multi-domain neural machine translation with word-level domain context discrimination. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 447–457, Brussels, Belgium. Association for Computational Linguistics.

A Complete Results

We present the complete results of our own implementation in Table 5, Table 6, Table 7, Table 8 and Table 9. The multilingual results for heuristic static strategies and MULTIDDS-S are obtained with the hyperparameters provided by Wang et al. (2020b).

B Hyperparameters for Optimization

Multilingual NMT For MULTIUAT, the NMT model is optimized with Adam (Kingma and Ba, 2015) with $\beta_1 = 0.9$ and $\beta_2 = 0.98$. The model is optimized for 40 epochs with the learning rate $\alpha = 5 \times 10^{-4}$ and the batch size of 9600 tokens. The learning rate increases linearly in the first $4K$ steps to the peak and then declines proportionally to the inverse square root of the number of steps. ψ is updated for every $2K$ steps with the learning rate 1×10^{-4} .

Multi-domain NMT The NMT model is optimized with Adam (Kingma and Ba, 2015) with $\beta_1 = 0.9$ and $\beta_2 = 0.98$. The model is optimized for 20 epochs with the learning rate $\alpha = 7 \times 10^{-4}$ and the batch size of $32K$ tokens. The learning rate increases linearly in the first $4K$ steps to the peak and then declines proportionally to the inverse square root of the number of steps. ψ for both MULTIDDS-S and MULTIUAT is updated for every $1K$ steps with the learning rate 1×10^{-4} . All the hyperparameters are identical among all the approaches in multi-domain setup.

	μ_{BLEU}	aze	bel	glg	slk	tur	rus	por	ces
PROP. ($\tau = 1$)	24.78	10.76	16.75	28.24	28.61	22.89	22.93	41.53	26.54
TEMP. ($\tau = 5$)	23.57	9.49	15.58	26.81	27.86	21.09	21.77	40.10	25.88
UNI. ($\tau = \infty$)	22.41	8.03	14.26	24.82	27.38	19.99	20.93	38.78	25.08
MULTIDDS-S	24.68	11.35	17.70	28.36	29.08	21.62	22.15	40.68	26.50
MULTIUAT									
+ PRETP	26.30	12.67†	19.28†	29.31†	30.57†	24.10†	23.75†	42.85†	27.89†
+ EXPTP	26.36	12.61†	18.81†	28.97	30.98†	24.56†	23.92†	42.83†	28.22†
+ VARTP	26.39	12.84†	19.79†	29.43†	30.63†	24.12†	23.68†	42.88†	27.77†
+ COMEV	26.26	13.06†	19.40†	29.31†	30.47†	23.81†	23.57†	42.77†	27.71†
+ ENTSENT	26.34	12.38†	19.75†	29.53†	30.40†	24.16†	23.86†	42.91†	27.69†
+ ENTEOS	26.28	12.72†	19.22†	29.40†	30.43†	24.02†	23.77†	42.68†	28.03†

Table 5: The complete results for M2O-related setting in multilingual NMT. Best results are highlighted in **bold**. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ given by Koehn (2004).

	μ_{BLEU}	bos	mar	hin	mkd	ell	bul	fra	kor
PROP. ($\tau = 1$)	26.28	21.87	9.92	21.51	31.12	35.21	36.02	37.75	16.83
TEMP. ($\tau = 5$)	25.77	23.18	9.72	21.01	30.61	34.29	34.98	36.49	15.90
UNI. ($\tau = \infty$)	24.97	21.48	9.27	20.25	30.08	33.70	34.11	35.54	15.30
MULTIDDS-S	26.67	24.64	10.45	22.09	32.31	35.13	35.50	37.06	16.16
MULTIUAT									
+ PRETP	27.82	26.05†	11.08	23.72†	32.39	35.92†	37.30†	38.71†	17.37†
+ EXPTP	27.74	24.96	11.28	23.60†	32.84	35.92†	37.08†	38.77†	17.47†
+ VARTP	27.83	25.45†	10.84	23.42†	32.89	36.28†	37.38†	38.89†	17.47†
+ COMEV	27.75	25.25	11.04	23.54†	33.00†	35.97†	37.21†	38.53†	17.42†
+ ENTSENT	27.71	24.95	11.25	23.30†	32.84	36.02†	37.22†	38.70†	17.42†
+ ENTEOS	27.82	25.90†	11.18	23.06†	33.36†	35.90†	37.26†	38.50†	17.38†

Table 6: The complete results for M2O-diverse setting in multilingual NMT. Best results are highlighted in **bold**. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ given by Koehn (2004).

	μ_{BLEU}	aze	bel	glg	slk	tur	rus	por	ces
PROP. ($\tau = 1$)	15.32	4.07	4.29	16.26	17.17	12.37	16.70	35.25	16.41
TEMP. ($\tau = 5$)	16.44	6.66	11.12	21.56	18.51	10.67	14.57	32.24	16.18
UNI. ($\tau = \infty$)	15.26	5.91	10.02	20.64	17.55	9.45	13.11	30.32	15.07
MULTIDDS-S	16.64	6.93	12.01	22.51	18.47	10.69	14.58	32.04	15.85
MULTIUAT									
+ PRETP	18.36	6.37	11.23	22.60	21.56†	13.41†	17.15†	35.88†	18.66†
+ EXPTP	18.64	6.95	12.25	23.50†	21.51†	13.37†	17.34†	35.86†	18.30†
+ VARTP	18.58	6.80	12.08	23.69†	21.08†	13.64†	17.31†	35.95†	18.06†
+ COMEV	18.58	7.10	12.04	23.43†	21.20†	13.32†	17.17†	35.90†	18.49†
+ ENTSENT	18.63	6.80	12.39	23.52†	21.37†	13.54†	17.31†	35.86†	18.22†
+ ENTEOS	18.56	6.41	11.95	23.05	21.70†	13.39†	17.17†	36.09†	18.68†

Table 7: The complete results for O2M-related setting in multilingual NMT. Best results are highlighted in **bold**. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ given by Koehn (2004).

	μ_{BLEU}	bos	mar	hin	mkd	ell	bul	fra	kor
PROP. ($\tau = 1$)	16.79	6.48	3.58	10.67	15.09	26.73	29.05	33.32	9.39
TEMP. ($\tau = 5$)	18.18	14.61	4.94	14.65	20.56	24.66	27.39	29.92	8.69
UNI. ($\tau = \infty$)	17.25	14.62	4.88	13.99	20.31	23.61	25.05	27.48	8.03
MULTIDDS-S	18.04	15.43	4.85	14.36	20.80	24.63	26.77	29.93	7.57
MULTIUAT									
+ PRETP	19.57	14.88	4.71	14.66	22.36†	27.21†	30.04†	33.89†	8.83†
+ EXPTP	19.59	15.11	4.86	14.64	22.06†	26.87†	30.22†	34.16†	8.81†
+ VARTP	19.62	15.35	4.67	14.52	22.13†	26.45†	30.43†	34.36†	9.01†
+ COMEV	19.67	14.68	4.72	14.73	23.12†	27.10†	30.23†	34.14†	8.62†
+ ENTSENT	19.68	14.51	4.69	14.74	22.70†	26.72†	30.52†	34.45†	9.09†
+ ENTEOS	19.76	14.59	4.83	14.63	23.08†	27.05†	30.64†	34.61†	8.68†

Table 8: The complete results for O2M-diverse setting in multilingual NMT. Best results are highlighted in **bold**. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ given by [Koehn \(2004\)](#).

	ID					OOD		
	μ_{BLEU}	WMT	Tanzil	EMEA	KDE	μ_{BLEU}	QED	TED
PROP. ($\tau = 1$)	35.69	24.76	40.23	45.56	32.19	25.15	22.71	27.59
TEMP. ($\tau = 5$)	36.92	21.63	38.49	51.53	36.03	23.46	21.13	25.78
UNI. ($\tau = \infty$)	36.85	20.17	38.89	52.17	36.16	22.22	19.85	24.58
MULTIDDS-S	36.42	21.76	37.41	51.92	34.60	22.74	20.30	25.18
MULTIUAT								
+ PRETP	37.50	22.21	39.61†	52.66	35.53†	23.78	21.20†	26.35†
+ EXPTP	37.29	22.22	39.37†	51.62†	35.94†	23.89	21.31†	26.46†
+ VARTP	35.48	23.08†	39.11†	46.91	32.80	23.75	21.21†	26.29†
+ COMBTP	37.25	21.73	38.51†	52.83†	35.91†	23.36	20.98†	25.74
+ ENTSENT	37.27	22.64†	38.64†	51.65	36.14†	23.86	21.29†	26.42†
+ ENTEOS	37.44	22.14	38.38†	52.81†	36.44†	23.80	21.34†	26.26†

Table 9: The complete results for multi-domain NMT. Best results are highlighted in **bold**. † indicates the improvement for the corresponding MULTIUAT result against MULTIDDS-S result is statistically significant at $p < 0.05$ given by [Koehn \(2004\)](#).