

Meta-Learning for Effective Multi-task and Multilingual Modelling

Ishan Tarunesh¹ Sushil Khyalia¹ Vishwajeet Kumar²
Ganesh Ramakrishnan¹ Preethi Jyothi¹

¹ Indian Institute of Technology Bombay

² IBM India Research Lab

{ishan, sushil, ganesh, pjyothi}@cse.iitb.ac.in
vishk024@in.ibm.com

Abstract

Natural language processing (NLP) tasks (e.g. question-answering in English) benefit from knowledge of other tasks (e.g., named entity recognition in English) and knowledge of other languages (e.g., question-answering in Spanish). Such shared representations are typically learned in isolation, either across tasks or across languages. In this work, we propose a meta-learning approach to learn the interactions between both tasks and languages. We also investigate the role of different sampling strategies used during meta-learning. We present experiments on five different tasks and six different languages from the XTREME multilingual benchmark dataset (Hu et al., 2020). Our meta-learned model clearly improves in performance compared to competitive baseline models that also include multi-task baselines. We also present zero-shot evaluations on unseen target languages to demonstrate the utility of our proposed model.

1 Introduction

Multi-task and multilingual learning are both problems of long standing interest in natural language processing. Leveraging data from multiple tasks and/or additional languages to benefit a target task is of great appeal, especially when the target task has limited resources. When it comes to multiple tasks, it is well-known from prior work on multi-task learning (Liu et al., 2019b; Kendall et al., 2018; Liu et al., 2019a; Yang and Hospedales, 2017) that jointly learning a model across tasks can benefit the tasks mutually. For multiple languages, the ability of deep learning models to learn effective embeddings has led to their use for joint learning of models across languages (Conneau et al., 2020; Conneau and Lample, 2019; Artetxe and Schwenk, 2019); learning cross-lingual embeddings to aid languages in limited resource settings is of growing interest (Kumar et al., 2019; Wang et al., 2017;

Adams et al., 2017). Let us say we had access to M tasks across N different languages - *c.f.* Table 1 that outlines such a matrix of tasks and languages from the XTREME benchmark (Hu et al., 2020). How do we perform effective joint learning across tasks and languages? Are there specific tasks or languages that need to be sampled more frequently for effective joint training? Can such sampling strategies be learned from the data?

In this work, we adopt a meta-learning approach for efficiently learning parameters in a shared parameter space across multiple tasks and multiple languages. Our chosen tasks are question answering (QA), natural language inference (NLI), paraphrase identification (PA), part-of-speech tagging (POS) and named entity recognition (NER). The tasks were chosen to enable us to employ a gamut of different types of language representations needed to tackle problems in NLP. In Figure 1, we illustrate the different types of representations by drawing inspiration from the Vauquois Triangle (Vauquois, 1968), well-known for machine translation, and situating our chosen tasks within such a triangle. Here we see that POS and NER are relatively ‘shallower’ analysis tasks that are token-centric, while QA, NLI and PA are ‘deeper’ analysis tasks that would require deeper semantic representations. This representation suggests a strategy for effective parameter sharing. For the deeper tasks, the same task in different languages could have representations that are closer and hence benefit each other, while for the shallower tasks, keeping the language unchanged and exploring different tasks might be more beneficial. Interestingly, this is exactly what we find with our meta-learned model and is borne out in our experimental results. We also find that as the model progressively learns, the meta-learning based models for the tasks requiring deeper semantic analysis benefit more from joint learning compared to the shallower tasks.

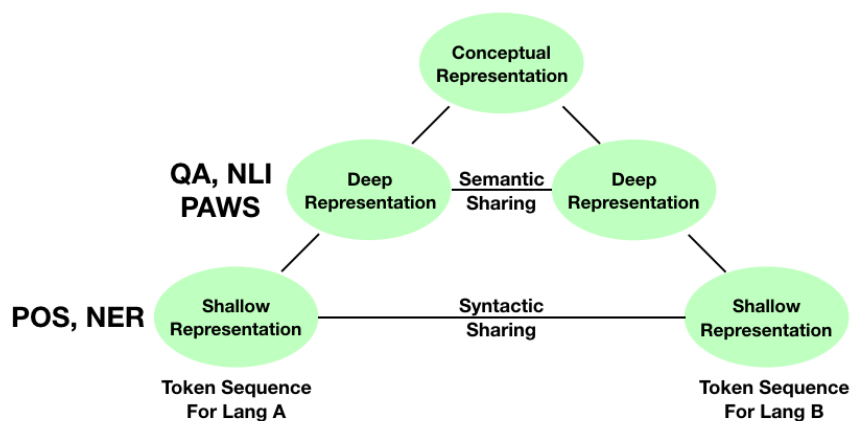


Figure 1: Illustration derived from Vauquois Triangle to linguistically motivate our setting. POS and NER being lower down in the representations (and are thus ‘shallower’) are further away from the same task in another language. QA, XNLI and PAWS being higher up in the representations (and are thus ‘deeper’) are closer to the same task in another language.

With access to multiple tasks and languages during training, the question of how to sample effectively from different tasks and languages also becomes important to consider. We investigate different sampling strategies, including a parameterized sampling strategy, to assess the influence of sampling across tasks and languages on our meta-learned model.

Our main contributions in this work are three-fold:

- We present a meta-learning approach that enables effective sharing of parameters across multiple tasks and multiple languages. This is the first work, to our knowledge, to explore the interplay between multiple tasks at different levels of abstraction and multiple languages using meta-learning. We show results on the recently-released XTREME benchmark and observe consistent improvements across different tasks and languages using our model. We also offer rules of thumb for effective meta-learning that could hold in larger settings involving additional tasks and languages.
- We investigate different sampling strategies that can be incorporated within our meta-learning approach and examine their benefits.
- We evaluate our meta-learned model in zero-shot settings for every task on target languages that never appear during training and show its superiority compared to competitive zero-shot baselines.

2 Related Work

We summarize three threads of related research that look at the transferability in models across different tasks and different languages: multi-task learning, meta-learning and data sampling strategies for both multi-task learning and meta-learning. Multi-task learning (Caruana, 1993) has proven to be highly effective for transfer learning in a variety of NLP applications such as question answering, neural machine translation, *etc.* (McCann et al., 2018; Hashimoto et al., 2017; Chen et al., 2018; Kipewasser and Ballesteros, 2018). Some multi-task learning approaches (Jawanpuria et al., 2015) have attempted to identify clusters (or groups) of related tasks based on end-to-end convex optimization formulations. Meta-learning algorithms (Nichol et al., 2018) are highly effective for fast adaptation and have recently been shown to be beneficial for several machine learning tasks (Santoro et al., 2016; Finn et al., 2017). Gu et al. (2018) use a meta-learning algorithm for machine translation to leverage information from high-resource languages. Dou et al. (2019) investigate multiple model agnostic meta-learning algorithms for low-resource natural language understanding on the GLUE (Wang et al., 2018) benchmark.

Data sampling strategies for multi-task learning and meta-learning form the third thread of related work. A good sampling strategy has to account for the imbalance in dataset sizes across tasks/languages and the similarity between tasks/languages. A simple heuristic-based solution to address the issue of data imbalance is to assign more weight to low-resource tasks or lan-

guages (Aharoni et al., 2019). Arivazhagan et al. (2019) define a temperature parameter which controls how often one samples from low-resource tasks/languages. The MultiDDS algorithm, proposed by Wang et al. (2020b), actively learns a different set of parameters for sampling batches given a set of tasks such that the performance on a held-out set is maximized. We use a variant of MultiDDS as a sampling strategy in our meta-learned model. Nooralahzadeh et al. (2020) is most similar in spirit to our work in that they study a cross-lingual and cross-task meta-learning architecture but only focus on zero-shot and few-shot transfer for two natural language understanding tasks, NLI and QA. In contrast, we study many tasks in many languages, in conjunction with sampling strategies, and offer concrete insights on how best to guide the meta-learning process when multiple tasks are in the picture.

3 Methodology

Our setting is pivoted on a **grid** of tasks and languages (with some missing entries as shown in Table 1). Each row of the grid corresponds to a single **task**. A cell of the grid corresponds to a Task-Language pair which we refer to as a TL pair (**TLP**). We denote by $q_i = |\mathcal{D}_{train}^i| / (\sum_{k=1}^n |\mathcal{D}_{train}^k|)$, the fraction of the dataset size for the i^{th} TLP and by $P_{\mathcal{D}}(i)$, the probability of sampling a batch from the i^{th} TLP during meta training. The distribution over all TLPs, *viz.*, is a Multinomial (say \mathcal{M}) over $P_{\mathcal{D}}(i)$ s.

3.1 Our Meta-learning Approach

The goal in the standard meta learning setting is to obtain a model that generalizes well to new test/target tasks given some distribution over training tasks. This can be achieved using optimization-based meta-learning algorithms that modify the learning procedure in order to learn a good initialization of the parameters. This can serve as a useful starting point that can be further fine-tuned on various tasks. Finn et al. (2017) proposed a general optimization algorithm called Model Agnostic Meta Learning (MAML) that can be trained using gradient descent. MAML aims to minimize the following objective

$$\min_{\theta} \sum_{T_i \sim \mathcal{M}} \mathcal{L}_i \left(U_i^k(\theta) \right) \quad (1)$$

where \mathcal{M} is the Multinomial distribution over TLPs, \mathcal{L}_i is the loss and U_i^k a function that returns θ

after k gradient updates both calculated on batches sampled from T_i . Minimizing this objective using first order methods involves computing gradients of the form $\frac{\partial}{\partial \theta} U_i^k(\theta)$, leading to the expensive computation of second order derivatives. Nichol et al. (2018) proposed an alternative first-order meta-learning algorithm named ‘‘Reptile’’ with simple update rule:

$$\theta \leftarrow \theta + \beta \frac{1}{|\{T_i\}|} \sum_{T_i \sim \mathcal{M}} (\theta_i^{(k)} - \theta) \quad (2)$$

where $\theta_i^{(k)}$ is $U_i^k(\theta)$. Despite its simplicity, a recent study by Dou et al. (2019) showed that Reptile is atleast as effective as MAML in terms of performance. We therefore employed Reptile for meta learning in all our experiments.

Algorithm 1 Our Meta-learning Approach

Input: \mathcal{D}_{train} set of TLPs for meta training (Also \mathcal{D}_{dev} for parametrised sampling)

Sampling Strategy (Temperature / MultiDDS)

Output: The converged multi-task multilingual model parameters θ^*

- 1: **Initialize** $P_{\mathcal{D}}(i)$ depending on the sampling strategy
 - 2: **while** not converged **do**
 - 3: \triangleright Perform Reptile Updates
 - 4: Sample m TLPs T_1, T_2, \dots, T_m from \mathcal{M}
 - 5: **for** $i = 1, 2, \dots, m$ **do**
 - 6: $\theta_i^{(k)} \leftarrow U_i^k(\theta)$, denoting k gradient updates from θ on batches of TLP T_i
 - 7: **end for**
 - 8: $\theta \leftarrow \theta + \frac{\beta}{m} \sum_{i=1}^m (\theta_i^{(k)} - \theta)$
 - 9: **if** Sampling Strategy \leftarrow MultiDDS **then**
 - 10: **for** $\mathcal{D}_{train}^i \in \mathcal{D}_{train}$ **do**
 - 11: $R(i; \theta) \leftarrow \cos(g_{dev}, g_{train})$, g_{dev} is gradient on $\{\mathcal{D}_{dev}\}$ and g_{train} is gradient on \mathcal{D}_{train}^i
 - 12: **end for**
 - 13: \triangleright Update Sampling Probabilities
 - 14: $d_{\psi} \leftarrow \sum_{i=1}^n R(i; \theta) \cdot \nabla_{\psi} \log(P_{\mathcal{D}}(i; \psi))$
 - 15: $\psi \leftarrow \text{GradientUpdate}(\psi, d_{\psi})$
 - 16: **end if**
 - 17: **end while**
-

3.2 Selection and Sampling Strategies

3.2.1 Selection

The choice of TLPs in meta-learning plays a vital role in influencing the model performance, as we

will see in more detail in Section 5. Apart from the use of all TLPs across both tasks and languages during training, selecting all languages for a given task (Gu et al., 2018) and selecting all tasks for a given language (Dou et al., 2019) are two other logical choices. We refer to the last two settings as being `Task-Limited` and `Lang-Limited`, respectively.

3.2.2 Heuristic Sampling

Once the TLPs for meta training (denoted by \mathcal{D}) have been selected, we need to sample TLPs from \mathcal{M} . We investigate temperature-based heuristic sampling (Arivazhagan et al., 2019) which defines the probability of any dataset as a function of its size. $P_{\mathcal{D}}(i) = q_i^{1/\tau} / (\sum_{k=1}^n q_k^{1/\tau})$ where $P_{\mathcal{D}}(i)$ is the probability of the i^{th} TLP to be sampled and τ is the temperature parameter. $\tau = 1$ reduces to sampling TLPs proportional to their dataset sizes and $\tau \rightarrow \infty$ reduces to sampling TLPs uniformly.

3.2.3 Parameterized Sampling

The sampling strategy defined in Section 3.2.2 remains constant throughout meta training and only depends on dataset sizes. Wang et al. (2020b) proposed a parameterized sampling technique called MultiDDS that builds on Differential Data Selection (DDS) (Wang et al., 2020a) for weighing multiple datasets. The $P_{\mathcal{D}}(i)$ are parameterized using ψ_i as $P_{\mathcal{D}}(i) = e^{\psi_i} / \sum_j e^{\psi_j}$ with the initial value of ψ satisfying $P_{\mathcal{D}}(i) = q_i$. The optimization for ψ and θ is performed in an alternating manner (Colson et al., 2007)

$$\psi^* = \operatorname{argmin}_{\psi} J(\theta^*(\psi), \mathcal{D}_{dev}) \quad (3)$$

$$\theta^*(\psi) = \operatorname{argmin}_{\theta} E_{x,y \sim P(T;\psi)} [l(x, y; \theta)] \quad (4)$$

$J(\theta, \mathcal{D}_{dev})$ is the objective function which we want to minimize over development set(s). The reward function, $R(x, y; \theta_t)$, is defined as:

$$R(x, y; \theta_t) \approx \underbrace{\nabla J(\theta_t, \mathcal{D}_{dev})^T}_{g_{dev}} \cdot \underbrace{\nabla_{\theta} l(x, y; \theta_{t-1})}_{g_{train}} \quad (5)$$

$$\approx \cos(g_{dev}, g_{train}) \quad (6)$$

ψ 's are updated using the REINFORCE (Williams, 1992) algorithm.

$$\psi_{t+1} \leftarrow \psi_t + R(x, y; \theta_t) \cdot \nabla_{\psi} \log(P(x, y; \psi)) \quad (7)$$

The Reptile meta-learning algorithm (along with details of the parameterized sampling strategy) is outlined in Algorithm 1.

4 Experimental Setup

4.1 Evaluation Benchmark

The recently released XTREME dataset (Hu et al., 2020) is a multilingual multi-task benchmark consisting of classification, structured prediction, QA and retrieval tasks. Each constituent task has associated datasets in multiple languages. The sources of POS and NER datasets are Universal Dependency v2.5 treebank (Nivre et al., 2020) and WikiAnn (Pan et al., 2017) respectively, with ground-truth labels available for each language. Large-scale datasets for QA, NLI and PA were originally available only for English. The PAWS-X (Yang et al., 2019) dataset contains machine-translated training pairs and human-translated evaluation pairs for PA. The authors of XTREME train a custom-built translation system to obtain translated datasets for QA and NLI. For the NLI task, we train using MultiNLI (Williams et al., 2018) and evaluate on XNLI (Conneau et al., 2018). For the QA task, SQuAD 1.1 (Rajpurkar et al., 2016) was used for training and MLQA (Lewis et al., 2019) for evaluation.

Regarding evaluation metrics, for QA we report F1 scores and for the other four tasks (PA, NLI, POS, NER) we report accuracy scores.

4.2 Implementation Details

BERT (Devlin et al., 2019) models yield state-of-the-art performance for many NLP tasks. Since we are dealing with datasets in multiple languages, we build our meta learning models on mBERT (Pires et al., 2019; Wu and Dredze, 2019) base architecture, implemented by Wolf et al. (2020), with output layers specific to each task. In our experiments, we use the AdamW (Loshchilov and Hutter, 2017) optimizer to make gradient-based updates to the model's parameters using batches from a particular TLP (Alg. 1, Line 6). This optimizer is shared across all the TLPs. When performing the meta-step (Alg. 1, Line 8), we use vanilla stochastic gradient descent (SGD) (Robbins and Monro, 1951) updates. Similarly, in the case of parameterized sampling the weights are updated (Alg. 1, Line 15) using vanilla SGD.

Meta training involves sampling a set of m tasks, taking k gradient update steps from the initial parameter to arrive at $\theta_i^{(k)}$ for task T_i and finally updating θ using the Reptile update rule. For meta-

Task	en	hi	es	de	fr	zh
Natural Language Inference (NLI)	392K		392K	392K	392K	
Question Answering (QA)	88.0K	82.4K	81.8K	80.0K		
Part Of Speech (POS)	21.2K	13.3K	28.4K	166K		7.9K
Named Entity Recognition (NER)	20K	5K	20K	20K	20K	20K
Paraphrase Identification (PA)	49.4K		49.4K	49.4K	49.4K	49.4K

Table 1: Dataset matrix showing datasets that are available (green) from the XTREME Benchmark. The number of training instances are also mentioned for each available dataset.

models we fix learning rate = $3e-5$ and dropout probability = 0.1 (provided by XTREME for reproduction of baselines). Grid search was performed on $m \in \{4, 8, 16\}$, $k \in \{2, 3, 4, 5\}$ and $\beta \in \{0.1, 0.5, 1.0\}$ for All TLPs model ($\tau = 1$). The best setting ($m = 8$, $k = 3$, $\beta = 1.0$) was selected based on validation score (accuracy or F1) averaged over all TLPs. These hyperparameters were kept constant for all further experiments. Each meta-learning model is trained for 5 epochs. We then finetune the meta model individually on each TLP and evaluate the results. Finetuning parameters vary for different task and are mentioned in Appendix B.

4.3 Data Selection and Sampling Strategies

We experiment with three different configurations for the set of TLPs to be considered during meta-learning: (a) using all tasks for a given language (Lang-Limited) (b) using all languages for a given task (Task-Limited) and (c) using all tasks and all languages (All TLPs). Since the dataset size varies across tasks (as also across languages), we use temperature sampling within each setting for $\tau = 1, 2, 5$ and ∞ . (In Table 4 of the Appendix C in the supplementary material, we report results for different choices of TLP selection and different values of the temperature.)

With respect to the **Input** in Algorithm 1, there are two sets of TLPs that need to be selected for parameterized sampling: \mathcal{D}_{train} and \mathcal{D}_{dev} . In order to analyse the effect of the choice of task and language, we experiment with the following 4 settings -

(a) $\mathcal{D}_{train} = \text{Lang-Limited}$, $\mathcal{D}_{dev} = \text{Target TLP}$

(b) $\mathcal{D}_{train} = \text{Task-Limited}$, $\mathcal{D}_{dev} = \text{Target TLP}$

(c) $\mathcal{D}_{train} = \text{All TLPs}$, $\mathcal{D}_{dev} = \text{Lang-Limited}$

(d) $\mathcal{D}_{train} = \text{All TLPs}$, $\mathcal{D}_{dev} =$

Task-Limited.

The models (a), (b) are referred to as mDDS and (c), (d) are called mDDS-Lang and mDDS-Task respectively. Results for these 4 models are reported in Table 2 alongside temperature sampling for comparison.

4.4 Baselines

Our first baseline system for each TLP uses mBERT-based models trained on data specific to each TLP, which is either available as ground-truth or in a translated form. We follow the same hyperparameter settings as reported in XTREME. We also present three multi-task learning (MTL) baseline systems: task limited (Task-Limited), language limited (Lang-Limited), and the use of all TLPs during training (All TLPs MTL). During MTL training, we concatenate and shuffle the selected datasets. The model is trained for 5 epochs with a learning rate of $5e-5$. We refer the reader to Appendix A for more training details.

5 Results and Analysis

Table 2 presents all our main results comparing different data selection and sampling strategies used for meta-learning. Each column corresponds to a target TLP; the best-performing meta-learned models for each target TLP within each data selection setting have been highlighted in colour. (Light-to-dark gradation reflects improvements in performance.) From Table 2, we see that our meta-learned models outperform the baseline systems across all the TLPs corresponding to QA, NLI and PA. (POS and NER also mostly benefit from meta-learning, but the margins of improvement are much smaller compared to the other tasks given the already high baseline scores).

Task-Limited vs Lang-Limited models. For QA and NLI, we observe that the Task-Limited models are *always better* than

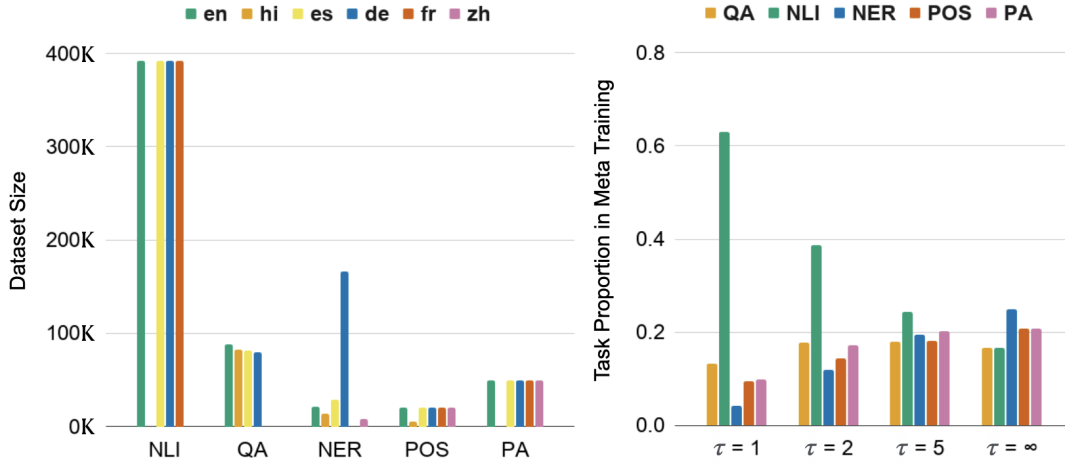


Figure 2: (a) Size of train dataset by language for each task (b) Proportion of dataset in meta training for different value of τ .

Model	SS	QA (F1)				NLI (Acc.)				PA (Acc.)				
		en	hi	es	de	en	es	de	fr	en	es	de	fr	zh
Baselines		79.94	59.94	65.83	63.17	81.39	78.37	76.82	77.30	92.35	89.75	87.45	89.61	83.32
Lang-Limited MTL		69.80	53.24	62.29	58.91	80.49	76.10	75.18	74.94	93.75	87.75	85.35	88.55	80.49
Task-Limited MTL		74.04	57.77	64.28	61.47	80.95	78.15	75.90	77.14	93.65	86.65	86.25	86.82	81.24
All TLPs MTL		63.22	42.94	54.05	51.61	80.05	76.48	74.86	76.18	93.50	90.30	88.45	89.71	82.66
Lang-Limited	Temp	-0.04	-0.24	-0.27	+0.07	+0.06	+0.39	+0.03	-0.70	+0.45	+0.05	+0.35	+0.40	-0.06
	mDDS	+0.07	-0.12	+0.06	+0.14	+0.02	-0.61	-0.80	-0.60	-0.25	-0.05	0.00	-0.30	-1.41
Task-Limited	Temp	+0.55	+0.43	+0.50	+0.40	+1.65	+1.12	+1.25	+0.79	+0.20	-0.15	-0.55	+0.85	-0.15
	mDDS	+0.21	+0.62	-0.67	+1.06	+1.32	+1.10	+1.39	+0.48	+0.50	-0.65	-0.35	+1.45	+1.06
All TLPs	Temp	+0.53	+0.47	+0.32	+0.47	+1.90	+1.22	+1.45	+0.95	+0.35	+0.45	+1.20	+1.05	+0.85
	mDDS-Lang	+0.08	+0.50	-1.57	+0.08	+0.76	+0.26	-0.10	+0.32	+0.25	+0.85	+0.75	+0.75	+1.11
	mDDS-Task	+0.18	+0.60	+0.11	+0.54	+1.50	+0.90	+0.72	+0.72	+0.10	+0.80	+1.27	+1.10	+1.16
Model	SS	NER (Acc.)					POS (Acc.)							
		en	hi	es	de	fr	zh	en	hi	es	de	zh		
Baselines		93.23	95.72	95.84	97.32	95.48	94.34	96.15	93.57	96.02	97.37	92.60		
Lang-Limited MTL		92.54	92.67	95.14	96.40	94.38	92.97	95.08	92.43	95.19	97.19	89.71		
Task-Limited MTL		93.51	93.94	95.77	97.09	95.27	93.72	95.70	93.34	95.73	97.35	92.52		
All TLPs MTL		92.28	91.95	94.90	96.18	94.38	92.53	94.70	91.89	95.10	97.03	89.92		
Lang-Limited	Temp	+0.60	+0.06	+0.09	+0.24	-0.09	-0.47	-0.06	-0.01	+0.10	+0.04	-0.17		
	mDDS	-0.21	-0.85	-0.20	-0.10	-0.57	-0.55	-0.27	-0.02	-0.19	-0.06	-0.37		
Task-Limited	Temp	+0.79	-0.46	0.00	-0.07	-0.18	-0.51	-0.22	-0.05	-0.21	+0.02	-0.09		
	mDDS	-0.10	-1.61	0.00	-0.16	-0.33	-0.69	-0.38	-0.02	-0.22	+0.05	-0.12		
All TLPs	Temp	-0.15	-0.70	+0.13	0.00	-0.16	-0.39	-0.22	-0.09	-0.21	+0.03	-0.16		
	mDDS-Lang	-0.16	-0.09	+0.11	-0.08	-0.14	-0.65	-0.21	-0.10	-0.11	+0.03	-0.17		
	mDDS-Task	-0.27	-0.42	+0.08	-0.14	-0.07	-0.58	-0.22	-0.14	-0.19	+0.02	-0.09		

Table 2: Main results comparing different data selection and sampling strategies. Sampling strategy, SS=Temp refers to the temperature-based sampling strategy and SS=mDDS refers to the multiDDS-based sampling strategy. mDDS-Lang and mDDS-Task refer to the use of a development set for multiDDS that contains all languages for a task and all tasks for a language, respectively. The best result among Baseline and three MTL models is highlighted using orange. For each column we present the difference (positive or negative) of the meta models from the best baseline (highlighted in orange) of that column

the Lang-Limited models. This is in line with our intuition that tasks like QA and NLI (which require deeper semantic representations) will benefit more by using data from different languages for the same task. We see the opposite seems to hold for POS and NER where the Lang-Limited models are *almost always better* than the Task-Limited models. With POS and NER being relatively shallower tasks, it makes

sense that they benefit more from language-specific training that relies on token embeddings shared across tasks.

Investigating Sampling Strategies. In Table 2, all the scores shown for the Temp sampling strategy are the best scores across four different values of T , $T = 1, 2, 5, \infty$. (The complete table is available in Appendix C in the supplementary material.) We also present comparisons with the

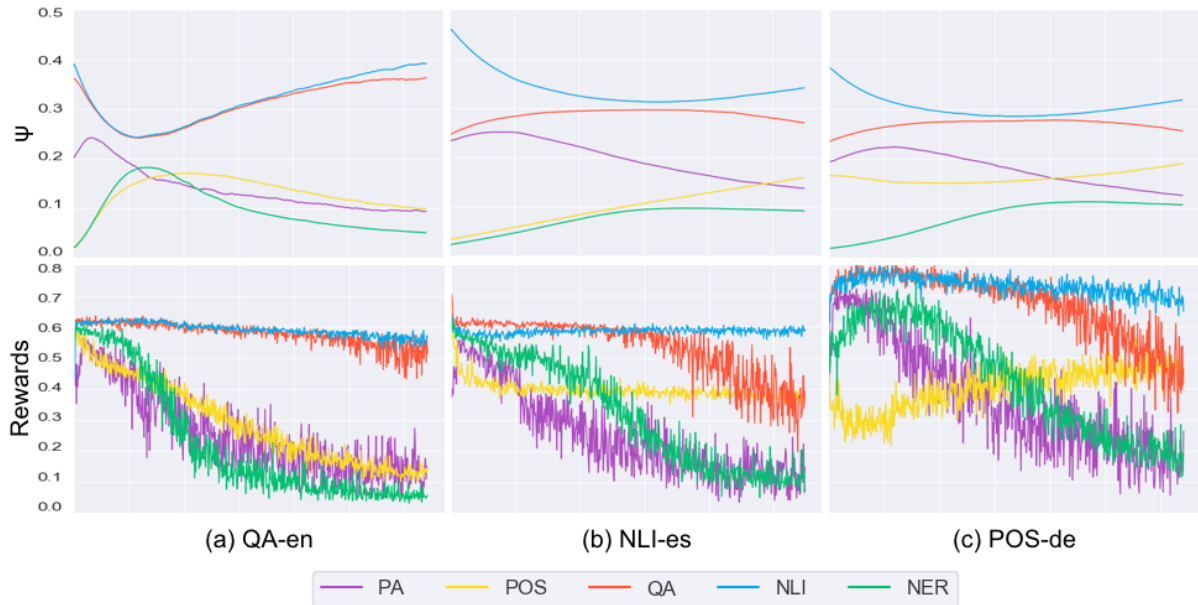


Figure 3: Evolution of ψ s and rewards as a function of training time for three Lang-Limited tasks evaluated on (a) QA-en (b) NLI-es and (c) POS-de.

mDDS, mDDS-Lang and mDDS-Task sampling strategies enforced within the Lang-Limited, Task-Limited and All TLPs models, respectively. For POS and NER, our best meta-learned models are mostly Lang-Limited with Temp sampling. It is intuitive that for these shallower tasks, mDDS does not offer any benefits from allowing to sample instances from other tasks.

To better understand the effects of mDDS sampling, Figure 3 shows plots of the rewards and sampling probabilities ψ 's computed as a function of training time for two deeper tasks - QA-en and NLI-es along with a shallower task - POS-de. We note that initially all the TLPs in any mDDS setting would start with similar rewards, thus lending ψ 's to converge towards the $T = \infty$ state. We highlight the following three observations:

- We find that the mDDS strategy does not help NLI at all. This is because the NLI task occupies the largest proportion across tasks at the start, as shown in Figure 2, and the proportion of NLI decreases substantially over time (since all tasks start with similar rewards at the beginning of meta training). Thus, for tasks that are over-represented in the meta-learning phase, temperature-based sampling is likely to be sufficient.
- We observe that the rewards for both QA and NLI are consistently high, irrespective of the

target TLP. This suggests that both QA and NLI are information-rich tasks and could benefit other tasks in meta-learning. This is also apparent from the accuracies for PA in Table 2, where all the best meta-learned models employ mDDS sampling.

- From the sampling probabilities for QA-en, we see that both QA and NLI are given almost equal weightage. However, from the F1 scores in Table 2, the best numbers for QA are in the Task-Limited setting which suggests that QA does not benefit from any other task. One explanation for this could be that the sequence length of inputs for NLI is 128 while the inputs for QA are of length 384, thus allowing lesser room for QA to be benefited by NLI.

Zero-shot Evaluations. Zero-shot evaluation is performed on languages that were not part of the training (henceforth, we refer to them as external languages). In the case of QA, NLI and PA we select all external language for which datasets were available in XTREME. For NER and POS, the number of external languages is close to 35 so we choose a subset of these to report the results. For evaluation, we compare models that are agnostic to the target language during meta training (Task-Limited, All TLPs and All TLPs mDDS-Task). Since Lang-Limited MTL is language specific and does not offer a competitive

Model	NER (Acc.)								POS (Acc.)							
	bn	et	fi	ja	mr	ta	te	ur	et	fi	ja	mr	ta	te	ur	
Task-Limited MTL	81.80	93.98	94.47	81.03	90.63	83.46	87.67	69.25	85.21	83.98	58.42	72.56	73.88	79.15	86.08	
All TLPs MTL	77.49	90.35	92.65	77.80	81.19	81.21	86.17	64.27	69.63	73.50	57.24	68.80	70.52	72.41	81.59	
Task-Limited	+1.91	+0.63	+0.16	+0.35	-0.67	+1.34	+0.63	+2.14	+2.94	+2.15	+0.83	+8.64	+2.34	+2.82	-0.30	
All TLPs	+0.62	+0.35	-0.11	+0.19	-0.92	+1.25	+0.43	+9.10	+2.56	+2.01	-1.42	+8.27	+1.24	+2.51	-0.16	
All TLPs mDDS-Task	-0.83	+0.09	-0.20	-1.34	-1.87	+0.49	+0.05	+3.62	+1.91	+1.08	-1.74	+8.64	+1.24	+1.88	-0.72	

Model	QA (F1)					NLI (Acc.)						PA (Acc.)		
	ar	vi	ar	bg	el	ru	sw	th	tr	ur	vi	ja	ko	
Task-Limited MTL	32.25	44.35	62.88	67.47	66.09	67.85	43.61	43.16	57.79	57.03	69.45	78.23	74.85	
All TLPs MTL	40.14	54.08	64.54	67.99	66.25	70.05	43.89	45.72	56.73	56.93	72.02	77.61	73.49	
Task-Limited	+8.14	+6.63	+4.35	+5.15	+4.62	+2.72	+8.51	+14.42	+6.79	+5.27	+1.3	+0.21	+1.81	
All TLPs	+5.24	+3.62	+4.41	+4.73	+4.79	+2.94	+11.44	+13.04	+7.05	+5.67	+1.24	+3.07	+4.57	
All TLPs mDDS-Task	+6.89	+6.29	+3.19	+4.33	+4.09	+2.38	+8.71	+13.16	+7.09	+4.41	+1.04	+2.81	+4.92	

Table 3: Results comparing Zero-shot evaluations for several external languages with competitive MTL baselines. The best MTL model is highlighted using orange. Rows for meta models show the difference (positive or negative) of the meta model result from the best MTL setting (orange) for that column

baseline when applied to an external language, we compare against Task-Limited MTL and All TLPs MTL that are more competitive.

An interesting observation from the zero shot results in Table 3 is that for every external language, on the ‘shallower’ NER and POS tasks, the Task-Limited variant of meta-learning performs better than both the variants of MTL, *viz.*, Task-Limited MTL and All TLPs MTL. In contrast, the ‘deeper’ tasks, *viz.*, QA, NLI and PA benefit more from the use of meta-learning using All TLPs setting, presumably because, as argued earlier, the deeper tasks tend to help each other more.

6 Conclusion

We present effective use of meta-learning for capturing task and language interactions in multi-task, multi-lingual settings. The effective use involves appropriate strategies for sampling tasks and languages as well as rough knowledge of the level of abstraction (deep vs. shallow representation) of that task. We present experiments on the XTREME multilingual benchmark dataset using five tasks and six languages. Our meta-learned model shows clear performance improvements over competitive baseline models. We observe that deeper tasks consistently benefit from meta-learning. Furthermore, shallower tasks benefit from deeper tasks when meta-learning is restricted to a single language. Finally, zero-shot evaluations for several external languages demonstrate the benefit of using meta-learning over two multi-task baselines while also reinforcing the linguistic insight that tasks requiring deeper representations tend to collaborate

better.

Acknowledgements

We thank anonymous reviewers for providing constructive feedback. We are grateful to IBM Research, India (specifically the IBM AI Horizon Networks - IIT Bombay initiative) for their support and sponsorship.

References

- Oliver Adams, Adam Makarucha, Graham Neubig, Steven Bird, and Trevor Cohn. 2017. Cross-lingual word embeddings for low-resource language modeling. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 937–947.
- Roei 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.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, M. Gatu Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *ArXiv*, abs/1907.05019.
- Mikel Artetxe and Holger Schwenk. 2019. Massively Multilingual Sentence Embeddings for Zero-Shot Cross-Lingual Transfer and Beyond. *Transactions of the ACL 2019*.
- Rich Caruana. 1993. Multitask learning: A knowledge-based source of inductive bias. In *ICML*.

- Junkun Chen, Xipeng Qiu, Pengfei Liu, and Xuanjing Huang. 2018. Meta multi-task learning for sequence modeling. In *AAAI*.
- Benoît Colson, Patrice Marcotte, and Gilles Savard. 2007. An overview of bilevel optimization. *Annals of operations research*, 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.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 7059–7069. Curran Associates, Inc.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Keyi Yu, and Antonios Anastasopoulos. 2019. Investigating meta-learning algorithms for low-resource natural language understanding tasks. 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 1192–1197.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135, International Convention Centre, Sydney, Australia. PMLR.
- Jiatao Gu, Yong Wang, Yun Chen, Victor OK Li, and Kyunghyun Cho. 2018. Meta-learning for low-resource neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. [A joint many-task model: Growing a neural network for multiple NLP tasks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1923–1933, Copenhagen, Denmark. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Pratik Jawanpuria, Jagarlapudi Saketha Nath, and Ganesh Ramakrishnan. 2015. [Generalized hierarchical kernel learning](#). *J. Mach. Learn. Res.*, 16:617–652.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491.
- Eliyahu Kiperwasser and Miguel Ballesteros. 2018. Scheduled multi-task learning: From syntax to translation. *Transactions of the Association for Computational Linguistics*, 6:225–240.
- Vishwajeet Kumar, Nitish Joshi, Arijit Mukherjee, Ganesh Ramakrishnan, and Preethi Jyothi. 2019. [Cross-lingual training for automatic question generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4863–4872, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Barlas Oğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. MLQA: Evaluating Cross-lingual Extractive Question Answering. *arXiv preprint arXiv:1910.07475*.
- Shikun Liu, Edward Johns, and Andrew J. Davison. 2019a. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. [Multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2017. Fixing weight decay regularization in adam. *ArXiv*, abs/1711.05101.

- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. The natural language decathlon: Multitask learning as question answering. *arXiv preprint arXiv:1806.08730*.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *ArXiv*, abs/1803.02999.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Farhad Nooralahzadeh, Giannis Bekoulis, Johannes Bjerva, and Isabelle Augenstein. 2020. [Zero-shot cross-lingual transfer with meta learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4547–4562, Online. Association for Computational Linguistics.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. Cross-lingual name tagging and linking for 282 languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407.
- Adam Santoro, Sergey Bartunov, Matthew M Botvinick, Daan Wierstra, and Timothy P. Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *ICML*.
- Bernard Vauquois. 1968. A survey of formal grammars and algorithms for recognition and transformation in mechanical translation. In *IFIP Congress*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Dingquan Wang, Nanyun Peng, and Kevin Duh. 2017. A multi-task learning approach to adapting bilingual word embeddings for cross-lingual named entity recognition. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 383–388.
- Xinyi Wang, Hieu Pham, Paul Michel, Antonios Anastasopoulos, Jaime Carbonell, and Graham Neubig. 2020a. [Optimizing data usage via differentiable rewards](#). In *Proceedings of the 37th International Conference on Machine Learning*, 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- 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.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. [PAWS-X: A cross-lingual adversarial dataset for paraphrase identification](#). 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 3687–3692, Hong Kong, China. Association for Computational Linguistics.

Yongxin Yang and Timothy M. Hospedales. 2017. [Deep multi-task representation learning: A tensor factorisation approach](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Appendices

Appendix A: Baseline Training Details

For QA learning rate is $3e-5$ and sequence length is 384 and the model is trained for 2 epochs. For PA, NLI, POS and NER the learning rate is $2e-5$ and sequence length is 128. NLI and PA models are trained for 5 epochs while POS and NER models are trained for 10 epochs. The choice of hyperparameters was kept constant across different languages for the same task.

Appendix B: Finetuning Details

For finetuning we kept the same number of epochs as the baseline of that task i.e 2 epochs for QA, 10 epochs for POS and NER, 5 epochs for NLI and PA. For QA we finetune with learning rate $3e-5$ and $3e-6$ and POS/NER we finetune with learning rate $2e-5$ and $2e-6$ and select the better of the two model. For PA and NLI the results for learning rate $2e-5$ were consistently worse compared to $2e-6$ so we just use $lr = 2e-6$ for PA and NLI.

Appendix C: Temperature Sampling

Model	T	QA (F1)				NLI (Acc.)				PA (Acc.)				
		en	hi	es	de	en	es	de	fr	en	es	de	fr	zh
Baselines		79.94	59.94	65.83	63.17	81.39	78.37	76.82	77.30	92.35	89.75	87.45	89.61	83.32
Lang-Limited MTL		69.80	53.24	62.29	58.91	80.49	76.10	75.18	74.94	93.75	87.75	85.35	88.55	80.49
Task-Limited MTL		74.04	57.77	64.28	61.47	80.95	78.15	75.90	77.14	93.65	86.65	86.25	86.82	81.24
All TLPs MTL		63.22	42.94	54.05	51.61	80.05	76.48	74.86	76.18	93.50	90.30	88.45	89.71	82.66
Lang-Limited	T = 1	79.49	59.42	64.67	63.04	81.13	78.76	76.23	76.51	93.85	89.15	87.83	89.63	82.56
	T = 2	78.81	59.68	65.10	63.24	80.87	77.56	76.85	76.60	93.85	90.15	87.70	89.41	83.10
	T = 5	79.90	58.74	65.56	62.12	81.19	78.17	76.10	76.56	93.65	90.35	88.60	90.11	83.20
	T = ∞	79.71	59.70	65.29	62.89	81.45	78.45	76.74	76.46	94.20	89.65	88.80	89.56	83.26
Task-Limited	T = 1	80.30	60.37	66.32	63.57	82.91	79.49	77.96	78.02	93.95	90.15	87.50	90.56	82.66
	T = 2	79.95	59.94	66.33	63.50	83.03	79.41	77.94	78.08	93.05	89.85	87.90	89.66	83.17
	T = 5	80.49	60.17	65.94	62.74	82.75	79.33	77.98	78.00	93.90	89.80	87.65	90.21	83.12
	T = ∞	79.77	59.86	66.01	62.96	83.03	79.39	78.07	78.09	93.60	89.75	87.75	89.61	82.42
All TLPs	T = 1	80.20	59.89	66.10	63.64	83.29	79.59	77.84	78.19	93.90	89.95	88.70	90.41	83.57
	T = 2	80.47	60.41	66.04	63.56	82.71	78.83	77.96	78.04	93.50	90.75	89.65	90.71	84.02
	T = 5	80.01	59.38	66.15	63.53	83.19	79.51	78.10	78.21	94.10	90.05	88.70	90.26	84.17
	T = ∞	80.27	59.82	64.41	63.08	83.27	79.43	78.27	78.25	94.05	90.75	88.70	90.76	83.42
Model	T	NER (Acc.)						POS (Acc.)						
		en	hi	es	de	fr	zh	en	hi	es	de	zh		
Baselines		93.23	95.72	95.84	97.32	95.48	94.34	96.15	93.57	96.02	97.37	92.60		
Lang-Limited MTL		92.54	92.67	95.14	96.40	94.38	92.97	95.08	92.43	95.19	97.19	89.71		
Task-Limited MTL		93.51	93.94	95.77	97.09	95.27	93.72	95.70	93.34	95.73	97.35	92.52		
All TLPs MTL		92.28	91.95	94.90	96.18	94.38	92.53	94.70	91.89	95.10	97.03	89.92		
Lang-Limited	T = 1	93.14	95.36	95.40	97.21	95.39	93.63	95.96	93.33	95.81	97.32	92.32		
	T = 2	93.24	94.76	95.80	97.56	95.07	93.53	95.87	93.53	95.93	97.39	92.40		
	T = 5	94.03	95.78	95.93	97.24	94.99	93.60	96.09	93.56	95.85	97.33	92.43		
	T = ∞	94.11	95.40	95.75	96.89	95.35	93.87	95.99	93.28	96.12	97.41	92.35		
Task-Limited	T = 1	94.30	95.26	95.82	97.25	95.26	93.62	95.93	93.36	95.81	97.31	92.38		
	T = 2	93.30	94.92	95.82	97.07	95.30	93.63	95.84	93.52	95.78	97.31	92.38		
	T = 5	93.29	95.02	95.73	96.98	95.19	93.56	95.92	93.34	95.75	97.39	92.43		
	T = ∞	93.37	94.70	95.84	96.95	95.20	93.83	95.77	93.33	95.76	97.33	92.51		
All TLPs	T = 1	93.14	93.63	95.91	97.30	95.32	93.53	95.90	93.35	95.76	97.36	92.43		
	T = 2	93.35	95.02	95.78	97.30	95.29	93.58	95.92	93.48	95.81	97.39	92.44		
	T = 5	93.36	94.51	95.93	97.26	95.28	93.95	95.92	93.35	95.78	97.40	92.42		
	T = ∞	93.35	94.95	95.97	97.32	95.28	93.63	95.93	93.31	95.80	97.30	92.43		

Table 4: Detailed results of temperature based heuristic sampling for different selections settings. The best result among Baseline and three MTL models is highlighted using orange. For each column we present the difference (positive or negative) of the meta models from the best baseline (highlighted in orange) of that column